

Cloudera Runtime 7.1.9

Release Notes

Date published: 2023-08-31

Date modified: 2026-06-13

CLOUdera

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

Legal Notice

© Cloudera Inc. 2026. All rights reserved.

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

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

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

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

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

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

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

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

Contents

Overview..... 9

7.1.9 SP1..... 9

What's new in Cloudera Runtime 7.1.9 SP1.....	9
Cloudera Runtime 7.1.9 SP1 Download URL.....	15
What's new in Platform Support.....	15
Cloudera Runtime 7.1.9 SP1 component versions.....	16
Cloudera Runtime 7.1.9.1073-8.....	16
Cloudera Runtime 7.1.9.1069-3.....	17
Cloudera Runtime 7.1.9.1066-8.....	19
Cloudera Runtime 7.1.9.1064-1.....	20
Cloudera Runtime 7.1.9.1059-4.....	21
Cloudera Runtime 7.1.9.1054-4.....	23
Cloudera Runtime 7.1.9.1045-5.....	24
Cloudera Runtime 7.1.9.1042-1.....	25
Cloudera Runtime 7.1.9.1037-2.....	27
Cloudera Runtime 7.1.9.1034-1.....	28
Cloudera Runtime 7.1.9.1032-3.....	29
Cloudera Runtime 7.1.9.1023-3.....	31
Cloudera Runtime 7.1.9.1015-6.....	32
Cloudera Runtime 7.1.9.1010-2.....	33
Cloudera Runtime 7.1.9.1005-2.....	35
Cloudera Runtime 7.1.9.1000-103.....	36
Using the Cloudera Runtime Maven repository 7.1.9 SP1.....	37
Runtime 7.1.9.1073-8.....	38
Runtime 7.1.9.1069-3.....	55
Runtime 7.1.9.1066-8.....	72
Runtime 7.1.9.1064-1.....	90
Runtime 7.1.9.1059-4.....	107
Runtime 7.1.9.1054-4.....	124
Runtime 7.1.9.1045-5.....	142
Runtime 7.1.9.1042-1.....	159
Runtime 7.1.9.1037-2.....	176
Runtime 7.1.9.1034-1.....	193
Runtime 7.1.9.1032-3.....	211
Runtime 7.1.9.1023-3.....	228
Runtime 7.1.9.1015-6.....	245
Runtime 7.1.9.1010-2.....	263
Runtime 7.1.9.1005-2.....	280
Runtime 7.1.9.1000-103.....	297
Fixed issues in Cloudera Runtime 7.1.9 SP1.....	305
Fixed Issues in Atlas.....	305
Fixed Issues in Arrow.....	307
Fixed Issues in Avro.....	307
Fixed issues in Calcite.....	308
Fixed issues in Cloud Connectors.....	308
Fixed issues in Cruise Control.....	308
Fixed Issues in Hadoop.....	308

Fixed Issues in HDFS.....	308
Fixed Issues in HBase.....	308
Fixed Issues in Hive.....	309
Fixed Issues in Hue.....	310
Fixed Issues in Impala.....	311
Fixed Issues in Iceberg.....	314
Fixed Issues in Kafka.....	315
Fixed Issues in Kerberos.....	315
Fixed Issues in Key HSM.....	316
Fixed Issues in Key Trustee KMS.....	316
Fixed Issues in Kudu.....	316
Fixed Issues in Knox.....	317
Fixed Issues in Livy.....	318
Fixed Issues in MapReduce.....	318
Fixed Issues in Navigator Encrypt.....	318
Fixed Issues in Oozie.....	319
Fixed issues in Ozone.....	319
Fixed Issues in ORC.....	321
Fixed Issues in Parquet.....	321
Fixed Issues in Phoenix.....	321
Fixed Issues in Ranger.....	321
Fixed Issues in Ranger KMS.....	324
Fixed Issues in Schema Registry.....	324
Fixed Issues in Solr.....	324
Fixed Issues in Spark.....	324
Fixed Issues in Spark Atlas Connector.....	324
Fixed Issues in Sqoop.....	325
Fixed Issues in Streams Replication Manager.....	325
Fixed Issues in Streams Messaging Manager.....	325
Fixed Issues in Tez.....	325
Fixed Issues in YARN and YARN Queue Manager.....	325
Fixed Issues in Zeppelin.....	326
Fixed Issues in Zookeeper.....	326
Known issues in Cloudera Runtime 7.1.9 SP1.....	326
Known Issues in Atlas.....	326
Known Issues in Avro.....	333
Known issues in Cloud Connectors.....	333
Known issues in Cruise Control.....	334
Known Issues in Hadoop.....	334
Known Issues in HBase.....	334
Known Issues in HDFS.....	336
Known Issues in Hive.....	338
Known Issues in Hue.....	343
Known Issues in Iceberg.....	349
Known Issues in Impala.....	349
Known Issues in Kafka.....	354
Known Issues in Kerberos.....	358
Known issues in Key Trustee Server.....	358
Known Issues in Knox.....	358
Known Issues in Kudu.....	360
Known Issues in Livy.....	360
Known Issues in Navigator Encrypt.....	360
Known Issues in Oozie.....	361
Known Issues in Ozone.....	362
Known Issues in Parquet.....	372
Known Issues in Phoenix.....	372

Known Issues in Queue Manager.....	373
Known Issues in Ranger.....	373
Known Issues in Ranger KMS.....	377
Known Issues in Schema Registry.....	377
Known Issues in Search Client.....	379
Known Issues in Solr.....	379
Known Issues in Spark.....	387
Known Issues in Streams Replication Manager.....	389
Known Issues for Sqoop.....	390
Known issues in Streams Messaging Manager.....	391
Known Issues in MapReduce, YARN and YARN Queue Manager.....	392
Known Issues in Zeppelin.....	399
Known Issues in ZooKeeper.....	399
Behavioral changes in Cloudera Runtime 7.1.9 SP1.....	400
Behavioral Changes in Ozone.....	400
CDP Private Cloud Base API Modifications and Removals.....	400
CDP 7.1.9 SP1 and 7.1.9 SP1 Cumulative Hotfix 1 Components with API differences.....	400
CDP 7.1.9 GA and 7.1.9 SP1 Components with API differences.....	401
CDP 7.1.9 Cumulative hotfix 7 and 7.1.9 SP1 Components with API differences.....	402
CDP 7.1.9 SP1 Components with API differences.....	404
Deprecation notices in Cloudera Runtime 7.1.9 SP1.....	404
Platform and OS.....	405
Ozone.....	405
Fixed Common Vulnerabilities and Exposures 7.1.9 SP1.....	405
Release Matrix.....	406
Cumulative hotfixes.....	408
Cumulative hotfix 15.....	408
Cumulative hotfix 14.....	463
Cumulative hotfix 13.....	518
Cumulative hotfix 12.....	575
Cumulative hotfix 11.....	631
Cumulative hotfix 10.....	689
Cumulative hotfix 9.....	741
Cumulative hotfix 8.....	794
Cumulative hotfix 7.....	847
Cumulative hotfix 6.....	902
Cumulative hotfix 5.....	957
Cumulative hotfix 4.....	965
Cumulative hotfix 3.....	974
Cumulative hotfix 2.....	983
Cumulative hotfix 1.....	991

7.1.9..... 1001

What's new in Cloudera Runtime 7.1.9.....	1001
Atlas.....	1004
Cruise Control.....	1005
HBase.....	1007
Hive.....	1007
Hue.....	1008
Iceberg.....	1010
Impala.....	1010
Kafka.....	1014
Kerberos.....	1018
Key Trustee Server.....	1019
Knox.....	1019

Kudu.....	1019
Livy.....	1020
Navigator Encrypt.....	1021
Oozie.....	1021
Ozone.....	1022
Phoenix.....	1023
Ranger.....	1024
Ranger KMS.....	1025
Schema Registry.....	1026
Solr.....	1027
Spark.....	1027
Sqoop.....	1028
SRM.....	1029
SMM.....	1030
YARN and YARN Queue Manager.....	1033
Zookeeper.....	1034
Unaffected Components in this release.....	1034
Cloudera Runtime 7.1.9 component versions.....	1035
Using the Cloudera Runtime Maven repository 7.1.9.....	1036
Runtime 7.1.9.0-387.....	1036
Runtime 7.1.9.2-10.....	1054
Runtime 7.1.9.3-4.....	1071
Runtime 7.1.9.4-4.....	1088
Runtime 7.1.9.6-3.....	1105
Runtime 7.1.9.7-2.....	1122
Runtime 7.1.9.9-1.....	1140
Runtime 7.1.9.14-2.....	1157
What's new in Platform Support.....	1174
Fixed issues in Cloudera Runtime 7.1.9.....	1174
Atlas.....	1174
Avro.....	1177
Cloud Connectors.....	1178
Cruise Control.....	1179
Hadoop.....	1179
HDFS.....	1179
HBase.....	1180
Hive.....	1180
Hue.....	1184
Impala.....	1184
Kafka.....	1186
Kudu.....	1187
Knox.....	1188
Livy.....	1189
Navigator Encrypt.....	1189
Oozie.....	1189
Ozone.....	1191
Parquet.....	1192
Phoenix.....	1192
Ranger.....	1195
Schema Registry.....	1201
Solr.....	1202
Spark.....	1203
Sqoop.....	1203
SRM.....	1204
SMM.....	1204
Tez.....	1206

YARN.....	1206
Zeppelin.....	1208
Zookeeper.....	1208
Known issues in Cloudera Runtime 7.1.9.....	1209
Atlas.....	1209
Avro.....	1215
Known issues in Cloud Connectors.....	1215
Known issues in Cruise Control.....	1215
Hadoop.....	1216
HBase.....	1216
HDFS.....	1218
Hive.....	1219
Hue.....	1226
Iceberg.....	1231
Impala.....	1232
Kafka.....	1237
Kerberos.....	1242
Key Trustee Server.....	1242
Knox.....	1242
Kudu.....	1243
Navigator Encrypt.....	1243
Oozie.....	1243
Ozone.....	1245
Parquet.....	1254
Phoenix.....	1254
Queue Manager.....	1255
Ranger.....	1255
Known Issues in Ranger KMS.....	1259
Schema Registry.....	1259
Search Client.....	1261
Solr.....	1261
Spark.....	1269
SRM.....	1270
Sqoop.....	1272
SMM.....	1273
YARN.....	1275
Zeppelin.....	1281
Zookeeper.....	1282
Behavioral changes in Cloudera Runtime 7.1.9.....	1283
Cruise Control.....	1283
Behavioral changes in Apache Hive.....	1283
Kafka.....	1286
Ozone.....	1286
Behavioral Changes in Apache Ranger.....	1286
Behavioral Changes in Schema Registry.....	1288
Solr.....	1289
Impala.....	1289
YARN.....	1289
CDP Private Cloud Base API Modifications and Removals.....	1289
CDP 7.1.7 SP2 and CDP 7.1.9 Components with API differences.....	1290
CDP 7.1.7 SP2 and CDP 7.1.9 Compatible Components.....	1312
CDP 7.1.8 CHF 12 and CDP 7.1.9 Components with API differences.....	1312
CDP 7.1.8 CHF 12 and CDP 7.1.9 Compatible Components.....	1352
CDP Private Cloud Base REST API Modifications and Removals.....	1352
CDP 7.1.7 SP2 and CDP 7.1.9 Components with REST API differences.....	1353
CDP 7.1.7 SP2 and CDP 7.1.9 Compatible Components.....	1355

CDP 7.1.8 CHF 12 and CDP 7.1.9 Components with REST API differences.....	1355
CDP 7.1.8 CHF 12 and CDP 7.1.9 Compatible Components.....	1357
Deprecation notices in Cloudera Runtime 7.1.9.....	1357
Platform and OS.....	1357
Deprecation Notices for Spark 2.....	1358
DAS.....	1358
Deprecation Notices for Zeppelin.....	1359
Deprecation Notices for Apache Oozie.....	1359
Deprecation Notices for Cruise Control.....	1359
Kafka.....	1359
Navigator Key Trustee Server.....	1360
Fixed Common Vulnerabilities and Exposures 7.1.9.....	1360
Cumulative hotfixes.....	1368
Cumulative hotfix 1.....	1368
Cumulative hotfix 2.....	1376
Cumulative hotfix 3.....	1387
Cumulative hotfix 4.....	1398
Cumulative hotfix 5.....	1406
Cumulative hotfix 6.....	1415
Cumulative hotfix 7.....	1422

Overview

You can review the Release Notes of Cloudera Runtime 7.1.9 for release-specific information related to new features and improvements, bug fixes, deprecated features and components, known issues, and changed features that can affect product behavior.

7.1.9 SP1

These are Release Notes for CDP Private Cloud Base 7.1.9 SP1.

Information for CDP Private Cloud Base 7.1.9 SP1 are found [here](#).

What's new in Cloudera Runtime 7.1.9 SP1

Understand the functionalities and improvements to features of components in Cloudera Runtime 7.1.9 SP1.

Atlas

Compression algorithm changed for Atlas HBase tables

The default compaction algorithm was changed from Gzip to Snappy for Atlas HBase tables. This can result in up to 50% lower compaction time. For more information, see the `compressionType` property in [HBase entities created in Atlas](#)

Atlas diagnostic bundle introduced

Atlas can provide you the necessary diagnostic data in a single GZ file inside the diagnostic zip package to help you and the Cloudera Support to troubleshoot problems. For more information on downloading and sending the compressed logs from the Cloudera Manager, see [Manually Triggering Collection and Transfer of Diagnostic Data to Cloudera](#).

Spark connector available

With the new Spark connector you can run Spark job for Atlas even when Kafka brokers are not available. For more information, see [Spark connector configuration in Apache Atlas](#).

Relationship Search is configurable

By default, Relationship Search is disabled by the new property to save time when starting or restarting Atlas. A new property is introduced in Atlas to configure Relationship Search. For more information, see [Using Relationship search](#).

Chinese, Japanese and Korean alphabet supported in searches

Chinese, Japanese and Korean characters are supported in the following search scenarios:

- Any type based searches including filter operators together with AND, OR filter operators
- Searches with multiple attribute filters, such as classification and description text
- Searches involving tags and operators
- Searches with custom attribute filters
- Searches including the "*" wildcard character

For more information on searching in Atlas, see [Using Free-text Search](#).

Atlas On-prem to On-prem Replication (Technical Preview)

Replicate governance metadata provided by Cloudera Atlas and data lineage using replication policies in Replication Manager.

This feature is available in this version of CDP but is not ready for production deployment.

Cloudera encourages you to explore this technical preview feature in non-production environments

and provide feedback on your experiences through the Cloudera Community Forums or through your account managers and field team contacts. For more information regarding limitations and unsupported features, see [Atlas replication policies \(technical preview\)](#).

Custom audit filter updates

The DELETE API call was updated for the single rule deletion. For more information, see [Using custom audit filters](#).

Audit aging enhancements and bug fixes

- Added note on use of regular expressions for configuring custom audit aging. For more information, see [Using custom audit aging](#).
- Added note that by default, Audit Aging is disabled through REST API. For more information, see [Using audit aging](#).
- Changed the default behavior of the `atlas.audit.sweep.out.entity.types` and `atlas.audit.sweep.out.action.types` properties. They are no longer applicable by default with sweep out option enabled. For more information, see [Using Sweep out configurations](#).

Hue

Added support for Python 3.10

Hue is now supported with Python 3.10 on Ubuntu 22, SLES 15 SP4, and SLES 15 SP5 operating systems. See [Installing Python 3](#).

Added support for PostgreSQL 16

Hue supports PostgreSQL 16 as its backend database.

Added support for Oracle 21c and 23ai

Hue supports Oracle 23ai (LTS, Latest) and Oracle 21c (Latest) as its backend database. For the supported `cx_Oracle` information and installation instructions, see [Using Oracle database with Hue](#).

Security improvements

- Upgraded the Interactive Python (IPython) command-line shell to version 8.10.0 to address CVE 2023-24816.
- Upgraded the Python library to version 3.10 to resolve the ReDoS (Regular expression Denial of Service) attacks and address CVE 2022-42969.
- Upgraded SQLParse to version 0.4.4 to prevent the ReDoS attacks and address CVE 2023-30608.

Added DEBUG-level logging in Hue server logs

The Hue servers logs now contain DEBUG-level information in addition to INFO-level information. Cloudera has fixed multiple logging issues, overall. Only INFO-level logging is enabled by default. However, you can turn on DEBUG-level logging from Cloudera Manager. See [Enabling debug logging for Hue server logs](#).

The requirement of installing pycopg2 package from the source on a FIPS-enabled cluster

If you use PostgreSQL as a backend database for Hue on a FIPS cluster on RHEL 8, you must install a version of the `psycopg2` package from the source to be at least 2.9.5 on all Hue hosts because the `psycopg2-binary` package uses its version of the `libssl` library file which does not support FIPS. See [Installing the pycopg2 Python package for PostgreSQL database on a FIPS cluster \(RHEL 8\)](#).

Upgrade considerations

If you are using MySQL as a backend database for Hue, then you must perform the upgrade tasks in the following order to avoid potential issues:

1. Install Python 3.10 on all Hue hosts.
2. Install the MySQL Client package.
3. Stop the Hue service.

4. Upgrade Cloudera Manager to CM 7.11.3 CHF7.
5. Start Cloudera Manager.
6. Upgrade CDP Private Cloud Base to 7.1.9 SP1.
7. Start the Hue service.

For detailed instructions about each step, see [Upgrading CDP Private Cloud Base to a higher version](#).



Note: You must start Hue only after upgrading Cloudera Manager, Python version, and CDP Private Cloud Base to 7.1.9 SP1.



Note: On CDP Private Cloud Base 7.1.9 SP1, Python 3.10 is not mandatory for using Hue with MySQL as its backend database.

Zero downtime upgrade (ZDU) considerations

If you are using the ZDU approach to upgrade to 7.1.9 SP1, then you must start Hue only after upgrading Cloudera Manager to 7.11.3 CHF 7, Python to version 3.10, and CDP Private Cloud Base to 7.1.9 SP1. Hue does not start after upgrading only Cloudera Manager and Python but not CDP.

Hive

Added support for Hive DatabaseType in JDBC storage handler

You can now connect to Apache Hive using JdbcStorageHandler to access the hive data source. For more information, see [Using JdbcStorageHandler to query external database](#).

HBase

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'
```

START_KEY	END_KEY	SIZE	REQ	LOCALITY	SERVER_NAME REGION_NAME
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
```

Kafka

Kafka connect fails to start on FIPS cluster

New properties are introduced for Kafka and Kafka Connect. These properties enable you to fine-tune the retry behavior of the Kafka and Kafka Connect Ranger plugins. Configuring these properties can help you avoid retry and timeout-related communication failures between the Kafka and Range services. The properties introduced are as follows:

- Ranger Kafka Plugin Policy Rest Client Retry Interval Ms
- Ranger Kafka Plugin Policy Rest Client Max Retry Attempts
- Ranger Kafka Connect Plugin Policy Rest Client Retry Interval Ms
- Ranger Kafka Connect Plugin Policy Rest Client Max Retry Attempts

For more information, see [Kafka Properties in Cloudera Runtime 7.1.9](#).

Kudu

Added support for Python 3.10

Kudu now supports Python 3.10. For more information regarding python setup see [Kudu Python client](#).

Impala

Added support for ZDU

Impala now supports zero-downtime upgrades (ZDU) in High Availability clusters. For details, see [Configuring components before starting ZDU](#).

Added support for Impala High Availability

Impala now supports High Availability (HA) by deploying pairs of StateStore and Catalog instances in primary/standby mode, ensuring continuous operation during failures. For details, see [Configuring Impala for High Availability](#).

Key HSM

CDPD-66512: Enabled FIPS for Luna 7 server

Added support for Key HSM with FIPS enabled Luna Server.

Key Trustee Server

KT-7452: Support added for RHEL 9

Added RHEL 9 support for KTS. You should manually install Python 2 and crontabs package in order to run keytrustee-server on RHEL 9.

KT-7581: Support added for RHEL 8.10

Added RHEL 8.10 support for KTS.

Navigator Encrypt

CDPD-61686: Support added for RHEL 9.2

Added RHEL 9.2 support for Navigator Encrypt.

KT-7440: Support added for Ubuntu 22

Added Ubuntu 22 support for NavEncrypt.

KT-7513: Navigator Encrypt generates keytab file when working with KMS

When using KMS as the key manager, Navigator Encrypt generates a keytab file for it to use to authenticate with Ranger KMS.

Ranger

Red Hat Enterprise Linux 8.10 support

Ranger KMS now supports Red Hat Enterprise Linux 8.10.

Ranger KMS

CDPD-66512: Enabled FIPS for Luna 7 server

Added support for Ranger KMS and Ranger KMS KTS with FIPS enabled Luna Server.

CDPD-67530: Support added for FIPS on JDK17

Ranger KMS now supports FIPS on JDK17.

CDPD-67536: Support added for Python 3.10

Ranger KMS now supports Python 3.10.

CDPD-68973/CDPD-68284: Support added for Oracle23c database

Ranger KMS supports Oracle23c database.

CDPD-69178: Support for ext libraries in higher versions of java

During Luna client installation, you need to copy Luna API and jar files to the JRE ext library. This ext directory does not exist in Java11 (or Java9+). Cloudera added support for ext libraries in higher versions of java.

Sqoop

Secure options to provide Hive password during a Sqoop import

When importing data into Hive using Sqoop and if LDAP authentication is enabled for Hive, the necessity to set the Hive password parameter directly in the command-line poses a potential vulnerability. Passwords provided in plaintext within command-line interfaces are susceptible to unauthorized access or interception, compromising sensitive credentials and, subsequently, the security of the entire data transfer process.

Learn about the secure options that you can use to provide the Hive password during Sqoop-Hive imports instead of the earlier way of providing the password as plaintext in the command-line interface.

CDS 3.3 Powered by Apache Spark

Apache Spark 3 integration with Schema Registry

Apache Spark 3 integrated with Schema Registry provides a library to leverage Schema Registry for managing Spark schemas and to serialize and/or de-serialize messages in Spark data sources and sinks. For more information, see [Apache Spark 3 integration with Schema Registry](#).

YARN Queue Manager

Support for YARN Queue Manager data Migration from H2 to PostgreSQL database

You can now migrate YARN Queue Manager from an H2 database to a PostgreSQL database post-installation or upgrade. All your existing data in the H2 database will be transferred to PostgreSQL. The YARN Queue Manager will establish connections using the PostgreSQL database.



Note: Migration supports only one-way transfer from H2 to PostgreSQL. Reverting to H2 after the migrating is not supported, and data migration from PostgreSQL back to H2 is not possible.

For more information, see [Migrating from H2 to PostgreSQL database in YARN Queue Manager](#).

YARN Queue Manager UI behaviour in mixed resource allocation mode

The mixed resource allocation mode in YARN is supported through safety valves. If you open the Queue Manager UI or access the Queue Manager APIs while this mode is active, the Queue Manager blocks access and informs you that mixed resource allocation is active and the Queue Manager is inaccessible until complete compatibility is achieved.

For more information, see [YARN Queue Manager UI behavior in mixed resource allocation mode](#).

Upgrade CDP 7.1.9 to CDP 7.1.9 SP1

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

Upgrade CDP 7.1.8 latest Cumulative hotfix to CDP 7.1.9 SP1

You can perform an In-place upgrade from CDP 7.1.8 latest Cumulative hotfix to CDP 7.1.9 SP1. For more information, see [CDP to CDP](#) documentation.

Upgrade CDP 7.1.7 SP3 to CDP 7.1.9 SP1

You can perform an In-place upgrade from CDP 7.1.7 SP3 to CDP 7.1.9 SP1. For more information, see [CDP to CDP](#) documentation.

Upgrade CDH 6.3.x to CDP 7.1.9 SP1

You can perform In-place upgrade from CDH6.3.x to CDP 7.1.9 SP1. For more information, see [CDH 6 to CDP](#) documentation.

Upgrade CDH 6.2.x to CDP 7.1.9 SP1

You can perform In-place upgrade from CDH6.2.x to CDP 7.1.9 SP1. For more information, see [CDH 6 to CDP](#) documentation.

Upgrade HDP 3 to CDP 7.1.9 SP1

You can perform In-place one-stage upgrade from HDP 3 to CDP 7.1.9 SP1 using CMA 3.3. For more information, see [HDP 3 to CDP](#) documentation.

Upgrade HDP 2 to CDP 7.1.9 SP1

You can perform In-place one-stage upgrade from HDP 2 to CDP 7.1.9 SP1 using CMA 3.3. However, you must first upgrade from HDP 2 to CDP 7.1.8 or CDP 7.1.7 SP2 and then upgrade to CDP 7.1.9 SP1. For more information, see [HDP 2 to CDP](#) documentation.

Rollback CDP CDP 7.1.9 SP1 to CDP 7.1.9

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

Rollback CDP 7.1.9 SP1 to CDP 7.1.8 latest Cumulative hotfix

You can downgrade or rollback an upgrade from CDP Private Cloud Base 7.1.9 SP1 to CDP 7.1.8 latest Cumulative hotfix. The rollback restores your CDP cluster to the state it was in before the upgrade, including the Kerberos and TLS/SSL configurations. For more information, see [Rollback CDP 7.1.9 SP1 to CDP 7.1.8 latest Cumulative hotfix](#) documentation.

Rollback CDP 7.1.9 SP1 to CDP 7.1.7 SP3

You can downgrade or rollback an upgrade from CDP Private Cloud Base 7.1.9 SP1 to CDP 7.1.7 SP3. The rollback restores your CDP cluster to the state it was in before the upgrade, including the Kerberos and TLS/SSL configurations. For more information, see [Rollback CDP 7.1.9 SP1 to CDP 7.1.7 SP3](#) documentation.

Rollback CDP 7.1.9 SP1 to CDH 6.3.x

You can downgrade or rollback an upgrade from CDP Private Cloud Base 7.1.9 SP1 to CDH 6.3.x. The rollback restores your CDH cluster to the state it was in before the upgrade, including the Kerberos and TLS/SSL configurations. For more information, see [Rollback CDP 7.1.9 SP1 to CDH 6.3.x](#) documentation.

Rollback CDP 7.1.9 SP1 to CDH 6.2.x

You can downgrade or rollback an upgrade from CDP Private Cloud Base 7.1.9 SP1 to CDH 6.2.x. The rollback restores your CDH cluster to the state it was in before the upgrade, including the Kerberos and TLS/SSL configurations. For more information, see [Rollback CDP 7.1.9 SP1 to CDH 6.2.x](#) documentation.

Rollback CDP 7.1.9 SP1 to HDP 3 or HDP 2

You can downgrade or rollback an upgrade from CDP Private Cloud Base 7.1.9 SP1 to HDP 3 or HDP 2. The rollback restores your CDH cluster to the state it was in before the upgrade, including the Kerberos and TLS/SSL configurations. For more information, see [Rollback CDP 7.1.9 SP1 to HDP 3 or HDP 2](#) documentation.

Cloudera Runtime 7.1.9 SP1 Download URL

Know more about the CDP 7.1.9 SP1. This service pack was released on July 19, 2024.

Cloudera Runtime 7.1.9.1000 download URL:

Parcel Repository Location

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1000/parcels/
```

What's new in Platform Support

You must be aware of the platform support for the Cloudera Runtime 7.1.9 SP1 release.

Platform Support Enhancements

New OS support:

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

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

New DB Versions:

- PostgreSQL 17 (from 7.1.9 SP1 CHF 9 onwards)
- PostgreSQL 16
- PostgreSQL 15
- MariaDB 10.11

- Oracle 23c
- Oracle 21c
- Oracle 19c

New JDK version: Support for FIPS + OpenJDK 17 (From Cloudera Runtime 7.1.9 SP1 release onwards). For more information, see [JDK FIPS prerequisites for CDP](#).

Cloudera Runtime 7.1.9 SP1 component versions

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

The component version number has three parts, [****Apache component version number****].[*Runtime version number*].[*Runtime Build number*]. For example, if the listed Apache HBase component version number is 2.4.17.7.1.9.1000-103, 2.4.17 is the upstream Apache HBase component version, 7.1.9 is the Runtime version, and 103 is Runtime build. You can also view the component version numbers in Cloudera Manager.

Cloudera Runtime 7.1.9.1073-8

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 15 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1073-8
Apache Atlas	2.1.0.7.1.9.1073-8
Apache Avro	1.11.1.7.1.9.1073-8
Apache Calcite	1.19.0.7.1.9.1073-8
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1073-8
Apache HBase	2.4.17.7.1.9.1073-8
Apache Hive	3.1.3000.7.1.9.1073-8
Apache Iceberg	1.3.0.7.1.9.1073-8
Apache Impala	4.0.0.7.1.9.1073-8
Apache Kafka	3.4.1.7.1.9.1073-8
Apache Knox	1.3.0.7.1.9.1073-8
Apache Kudu	1.17.0.7.1.9.1073-8
Apache Livy	0.7.2.7.1.9.1073-8
Apache MapReduce	3.1.1.7.1.9.1073-8
Apache Ozone	1.4.0.7.1.9.1073-8
Apache Oozie	5.1.0.7.1.9.1073-8
Apache ORC	1.5.1.7.1.9.1073-8
Apache Parquet	1.10.99.7.1.9.1073-8
Apache Phoenix	5.1.1.7.1.9.1073-8
Apache Ranger	2.4.0.7.1.9.1073-8
Apache Solr	8.11.2.7.1.9.1073-8
Apache Spark 2.x	2.4.8.7.1.9.1073-8

Component	Version
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1073-8
Apache Tez	0.9.1.7.1.9.1073-8
Apache Zeppelin	0.8.2.7.1.9.1073-8
Apache ZooKeeper	3.8.1.7.1.9.1073-8

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1073-8
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1073-8
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1073-8
Search	1.0.0.7.1.9.1073-8
Schema Registry	0.10.0.7.1.9.1073-8
Streams Messaging Manager	2.3.0.7.1.9.1073-8
Streams Replication Manager	1.1.0.7.1.9.1073-8

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1073-8
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1073-8
Spark Atlas Connector	0.1.0.7.1.9.1073-8
Spark Schema Registry	1.1.0.7.1.9.1073-8

Cloudera Runtime 7.1.9.1069-3

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 14 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1069-3
Apache Atlas	2.1.0.7.1.9.1069-3
Apache Avro	1.11.1.7.1.9.1069-3
Apache Calcite	1.19.0.7.1.9.1069-3
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1069-3
Apache HBase	2.4.17.7.1.9.1069-3
Apache Hive	3.1.3000.7.1.9.1069-3
Apache Iceberg	1.3.0.7.1.9.1069-3

Component	Version
Apache Impala	4.0.0.7.1.9.1069-3
Apache Kafka	3.4.1.7.1.9.1069-3
Apache Knox	1.3.0.7.1.9.1069-3
Apache Kudu	1.17.0.7.1.9.1069-3
Apache Livy	0.7.2.7.1.9.1069-3
Apache MapReduce	3.1.1.7.1.9.1069-3
Apache Ozone	1.4.0.7.1.9.1069-3
Apache Oozie	5.1.0.7.1.9.1069-3
Apache ORC	1.5.1.7.1.9.1069-3
Apache Parquet	1.10.99.7.1.9.1069-3
Apache Phoenix	5.1.1.7.1.9.1069-3
Apache Ranger	2.4.0.7.1.9.1069-3
Apache Solr	8.11.2.7.1.9.1069-3
Apache Spark 2.x	2.4.8.7.1.9.1069-3
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1069-3
Apache Tez	0.9.1.7.1.9.1069-3
Apache Zeppelin	0.8.2.7.1.9.1069-3
Apache ZooKeeper	3.8.1.7.1.9.1069-3

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1069-3
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1069-3
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1069-3
Search	1.0.0.7.1.9.1069-3
Schema Registry	0.10.0.7.1.9.1069-3
Streams Messaging Manager	2.3.0.7.1.9.1069-3
Streams Replication Manager	1.1.0.7.1.9.1069-3

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1069-3
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1069-3
Spark Atlas Connector	0.1.0.7.1.9.1069-3
Spark Schema Registry	1.1.0.7.1.9.1069-3

Cloudera Runtime 7.1.9.1066-8

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 13 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1066-8
Apache Atlas	2.1.0.7.1.9.1066-8
Apache Avro	1.11.1.7.1.9.1066-8
Apache Calcite	1.19.0.7.1.9.1066-8
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1066-8
Apache HBase	2.4.17.7.1.9.1066-8
Apache Hive	3.1.3000.7.1.9.1066-8
Apache Iceberg	1.3.0.7.1.9.1066-8
Apache Impala	4.0.0.7.1.9.1066-8
Apache Kafka	3.4.1.7.1.9.1066-8
Apache Knox	1.3.0.7.1.9.1066-8
Apache Kudu	1.17.0.7.1.9.1066-8
Apache Livy	0.7.2.7.1.9.1066-8
Apache MapReduce	3.1.1.7.1.9.1066-8
Apache Ozone	1.4.0.7.1.9.1066-8
Apache Oozie	5.1.0.7.1.9.1066-8
Apache ORC	1.5.1.7.1.9.1066-8
Apache Parquet	1.10.99.7.1.9.1066-8
Apache Phoenix	5.1.1.7.1.9.1066-8
Apache Ranger	2.4.0.7.1.9.1066-8
Apache Solr	8.11.2.7.1.9.1066-8
Apache Spark 2.x	2.4.8.7.1.9.1066-8
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1066-8
Apache Tez	0.9.1.7.1.9.1066-8
Apache Zeppelin	0.8.2.7.1.9.1066-8
Apache ZooKeeper	3.8.1.7.1.9.1066-8

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1066-8
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1066-8
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1066-8

Component	Version
Search	1.0.0.7.1.9.1066-8
Schema Registry	0.10.0.7.1.9.1066-8
Streams Messaging Manager	2.3.0.7.1.9.1066-8
Streams Replication Manager	1.1.0.7.1.9.1066-8

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1066-8
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1066-8
Spark Atlas Connector	0.1.0.7.1.9.1066-8
Spark Schema Registry	1.1.0.7.1.9.1066-8

Cloudera Runtime 7.1.9.1064-1

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 12 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1064-1
Apache Atlas	2.1.0.7.1.9.1064-1
Apache Avro	1.11.1.7.1.9.1064-1
Apache Calcite	1.19.0.7.1.9.1064-1
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1064-1
Apache HBase	2.4.17.7.1.9.1064-1
Apache Hive	3.1.3000.7.1.9.1064-1
Apache Iceberg	1.3.0.7.1.9.1064-1
Apache Impala	4.0.0.7.1.9.1064-1
Apache Kafka	3.4.1.7.1.9.1064-1
Apache Knox	1.3.0.7.1.9.1064-1
Apache Kudu	1.17.0.7.1.9.1064-1
Apache Livy	0.7.2.7.1.9.1064-1
Apache MapReduce	3.1.1.7.1.9.1064-1
Apache Ozone	1.4.0.7.1.9.1064-1
Apache Oozie	5.1.0.7.1.9.1064-1
Apache ORC	1.5.1.7.1.9.1064-1
Apache Parquet	1.10.99.7.1.9.1064-1
Apache Phoenix	5.1.1.7.1.9.1064-1
Apache Ranger	2.4.0.7.1.9.1064-1

Component	Version
Apache Solr	8.11.2.7.1.9.1064-1
Apache Spark 2.x	2.4.8.7.1.9.1064-1
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1064-1
Apache Tez	0.9.1.7.1.9.1064-1
Apache Zeppelin	0.8.2.7.1.9.1064-1
Apache ZooKeeper	3.8.1.7.1.9.1064-1

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1064-1
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1064-1
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1064-1
Search	1.0.0.7.1.9.1064-1
Schema Registry	0.10.0.7.1.9.1064-1
Streams Messaging Manager	2.3.0.7.1.9.1064-1
Streams Replication Manager	1.1.0.7.1.9.1064-1

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1064-1
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1064-1
Spark Atlas Connector	0.1.0.7.1.9.1064-1
Spark Schema Registry	1.1.0.7.1.9.1064-1

Cloudera Runtime 7.1.9.1059-4

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 11 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1059-4
Apache Atlas	2.1.0.7.1.9.1059-4
Apache Avro	1.11.1.7.1.9.1059-4
Apache Calcite	1.19.0.7.1.9.1059-4
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1059-4
Apache HBase	2.4.17.7.1.9.1059-4

Component	Version
Apache Hive	3.1.3000.7.1.9.1059-4
Apache Iceberg	1.3.0.7.1.9.1059-4
Apache Impala	4.0.0.7.1.9.1059-4
Apache Kafka	3.4.1.7.1.9.1059-4
Apache Knox	1.3.0.7.1.9.1059-4
Apache Kudu	1.17.0.7.1.9.1059-4
Apache Livy	0.7.2.7.1.9.1059-4
Apache MapReduce	3.1.1.7.1.9.1059-4
Apache Ozone	1.4.0.7.1.9.1059-4
Apache Oozie	5.1.0.7.1.9.1059-4
Apache ORC	1.5.1.7.1.9.1059-4
Apache Parquet	1.10.99.7.1.9.1059-4
Apache Phoenix	5.1.1.7.1.9.1059-4
Apache Ranger	2.4.0.7.1.9.1059-4
Apache Solr	8.11.2.7.1.9.1059-4
Apache Spark 2.x	2.4.8.7.1.9.1059-4
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1059-4
Apache Tez	0.9.1.7.1.9.1059-4
Apache Zeppelin	0.8.2.7.1.9.1059-4
Apache ZooKeeper	3.8.1.7.1.9.1059-4

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1059-4
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1059-4
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1059-4
Search	1.0.0.7.1.9.1059-4
Schema Registry	0.10.0.7.1.9.1059-4
Streams Messaging Manager	2.3.0.7.1.9.1059-4
Streams Replication Manager	1.1.0.7.1.9.1059-4

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1059-4
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1059-4

Component	Version
Spark Atlas Connector	0.1.0.7.1.9.1059-4
Spark Schema Registry	1.1.0.7.1.9.1059-4

Cloudera Runtime 7.1.9.1054-4

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 10 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1054-4
Apache Atlas	2.1.0.7.1.9.1054-4
Apache Avro	1.11.1.7.1.9.1054-4
Apache Calcite	1.19.0.7.1.9.1054-4
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1054-4
Apache HBase	2.4.17.7.1.9.1054-4
Apache Hive	3.1.3000.7.1.9.1054-4
Apache Iceberg	1.3.0.7.1.9.1054-4
Apache Impala	4.0.0.7.1.9.1054-4
Apache Kafka	3.4.1.7.1.9.1054-4
Apache Knox	1.3.0.7.1.9.1054-4
Apache Kudu	1.17.0.7.1.9.1054-4
Apache Livy	0.7.2.7.1.9.1054-4
Apache MapReduce	3.1.1.7.1.9.1054-4
Apache Ozone	1.4.0.7.1.9.1054-4
Apache Oozie	5.1.0.7.1.9.1054-4
Apache ORC	1.5.1.7.1.9.1054-4
Apache Parquet	1.10.99.7.1.9.1054-4
Apache Phoenix	5.1.1.7.1.9.1054-4
Apache Ranger	2.4.0.7.1.9.1054-4
Apache Solr	8.11.2.7.1.9.1054-4
Apache Spark 2.x	2.4.8.7.1.9.1054-4
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1054-4
Apache Tez	0.9.1.7.1.9.1054-4
Apache Zeppelin	0.8.2.7.1.9.1054-4
Apache ZooKeeper	3.8.1.7.1.9.1054-4

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1054-4

Component	Version
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1054-4
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1054-4
Search	1.0.0.7.1.9.1054-4
Schema Registry	0.10.0.7.1.9.1054-4
Streams Messaging Manager	2.3.0.7.1.9.1054-4
Streams Replication Manager	1.1.0.7.1.9.1054-4

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1054-4
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1054-4
Spark Atlas Connector	0.1.0.7.1.9.1054-4
Spark Schema Registry	1.1.0.7.1.9.1054-4

Cloudera Runtime 7.1.9.1045-5

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 9 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1045-5
Apache Atlas	2.1.0.7.1.9.1045-5
Apache Avro	1.11.1.7.1.9.1045-5
Apache Calcite	1.19.0.7.1.9.1045-5
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1045-5
Apache HBase	2.4.17.7.1.9.1045-5
Apache Hive	3.1.3000.7.1.9.1045-5
Apache Iceberg	1.3.0.7.1.9.1045-5
Apache Impala	4.0.0.7.1.9.1045-5
Apache Kafka	3.4.1.7.1.9.1045-5
Apache Knox	1.3.0.7.1.9.1045-5
Apache Kudu	1.17.0.7.1.9.1045-5
Apache Livy	0.7.2.7.1.9.1045-5
Apache MapReduce	3.1.1.7.1.9.1045-5
Apache Ozone	1.4.0.7.1.9.1045-5
Apache Oozie	5.1.0.7.1.9.1045-5

Component	Version
Apache ORC	1.5.1.7.1.9.1045-5
Apache Parquet	1.10.99.7.1.9.1045-5
Apache Phoenix	5.1.1.7.1.9.1045-5
Apache Ranger	2.4.0.7.1.9.1045-5
Apache Solr	8.11.2.7.1.9.1045-5
Apache Spark 2.x	2.4.8.7.1.9.1045-5
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1045-5
Apache Tez	0.9.1.7.1.9.1045-5
Apache Zeppelin	0.8.2.7.1.9.1045-5
Apache ZooKeeper	3.8.1.7.1.9.1045-5

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1045-5
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1045-5
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1045-5
Search	1.0.0.7.1.9.1045-5
Schema Registry	0.10.0.7.1.9.1045-5
Streams Messaging Manager	2.3.0.7.1.9.1045-5
Streams Replication Manager	1.1.0.7.1.9.1045-5

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1045-5
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1045-5
Spark Atlas Connector	0.1.0.7.1.9.1045-5
Spark Schema Registry	1.1.0.7.1.9.1045-5

Cloudera Runtime 7.1.9.1042-1

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 8 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1042-1
Apache Atlas	2.1.0.7.1.9.1042-1

Component	Version
Apache Avro	1.11.1.7.1.9.1042-1
Apache Calcite	1.19.0.7.1.9.1042-1
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1042-1
Apache HBase	2.4.17.7.1.9.1042-1
Apache Hive	3.1.3000.7.1.9.1042-1
Apache Iceberg	1.3.0.7.1.9.1042-1
Apache Impala	4.0.0.7.1.9.1042-1
Apache Kafka	3.4.1.7.1.9.1042-1
Apache Knox	1.3.0.7.1.9.1042-1
Apache Kudu	1.17.0.7.1.9.1042-1
Apache Livy	0.7.2.7.1.9.1042-1
Apache MapReduce	3.1.1.7.1.9.1042-1
Apache Ozone	1.4.0.7.1.9.1042-1
Apache Oozie	5.1.0.7.1.9.1042-1
Apache ORC	1.5.1.7.1.9.1042-1
Apache Parquet	1.10.99.7.1.9.1042-1
Apache Phoenix	5.1.1.7.1.9.1042-1
Apache Ranger	2.4.0.7.1.9.1042-1
Apache Solr	8.11.2.7.1.9.1042-1
Apache Spark 2.x	2.4.8.7.1.9.1042-1
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1042-1
Apache Tez	0.9.1.7.1.9.1042-1
Apache Zeppelin	0.8.2.7.1.9.1042-1
Apache ZooKeeper	3.8.1.7.1.9.1042-1

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1042-1
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1042-1
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1042-1
Search	1.0.0.7.1.9.1042-1
Schema Registry	0.10.0.7.1.9.1042-1
Streams Messaging Manager	2.3.0.7.1.9.1042-1
Streams Replication Manager	1.1.0.7.1.9.1042-1

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1042-1
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1042-1
Spark Atlas Connector	0.1.0.7.1.9.1042-1
Spark Schema Registry	1.1.0.7.1.9.1042-1

Cloudera Runtime 7.1.9.1037-2

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 7 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1037-2
Apache Atlas	2.1.0.7.1.9.1037-2
Apache Avro	1.11.1.7.1.9.1037-2
Apache Calcite	1.19.0.7.1.9.1037-2
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1037-2
Apache HBase	2.4.17.7.1.9.1037-2
Apache Hive	3.1.3000.7.1.9.1037-2
Apache Iceberg	1.3.0.7.1.9.1037-2
Apache Impala	4.0.0.7.1.9.1037-2
Apache Kafka	3.4.1.7.1.9.1037-2
Apache Knox	1.3.0.7.1.9.1037-2
Apache Kudu	1.17.0.7.1.9.1037-2
Apache Livy	0.7.2.7.1.9.1037-2
Apache MapReduce	3.1.1.7.1.9.1037-2
Apache Ozone	1.4.0.7.1.9.1037-2
Apache Oozie	5.1.0.7.1.9.1037-2
Apache ORC	1.5.1.7.1.9.1037-2
Apache Parquet	1.10.99.7.1.9.1037-2
Apache Phoenix	5.1.1.7.1.9.1037-2
Apache Ranger	2.4.0.7.1.9.1037-2
Apache Solr	8.11.2.7.1.9.1037-2
Apache Spark 2.x	2.4.8.7.1.9.1037-2
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1037-2
Apache Tez	0.9.1.7.1.9.1037-2
Apache Zeppelin	0.8.2.7.1.9.1037-2
Apache ZooKeeper	3.8.1.7.1.9.1037-2

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1037-2
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1037-2
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1037-2
Search	1.0.0.7.1.9.1037-2
Schema Registry	0.10.0.7.1.9.1037-2
Streams Messaging Manager	2.3.0.7.1.9.1037-2
Streams Replication Manager	1.1.0.7.1.9.1037-2

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1037-2
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1037-2
Spark Atlas Connector	0.1.0.7.1.9.1037-2
Spark Schema Registry	1.1.0.7.1.9.1037-2

Cloudera Runtime 7.1.9.1034-1

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 6 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1034-1
Apache Atlas	2.1.0.7.1.9.1034-1
Apache Avro	1.11.1.7.1.9.1034-1
Apache Calcite	1.19.0.7.1.9.1034-1
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1034-1
Apache HBase	2.4.17.7.1.9.1034-1
Apache Hive	3.1.3000.7.1.9.1034-1
Apache Iceberg	1.3.0.7.1.9.1034-1
Apache Impala	4.0.0.7.1.9.1034-1
Apache Kafka	3.4.1.7.1.9.1034-1
Apache Knox	1.3.0.7.1.9.1034-1
Apache Kudu	1.17.0.7.1.9.1034-1
Apache Livy	0.7.2.7.1.9.1034-1
Apache MapReduce	3.1.1.7.1.9.1034-1

Component	Version
Apache Ozone	1.4.0.7.1.9.1034-1
Apache Oozie	5.1.0.7.1.9.1034-1
Apache ORC	1.5.1.7.1.9.1034-1
Apache Parquet	1.10.99.7.1.9.1034-1
Apache Phoenix	5.1.1.7.1.9.1034-1
Apache Ranger	2.4.0.7.1.9.1034-1
Apache Solr	8.11.2.7.1.9.1034-1
Apache Spark 2.x	2.4.8.7.1.9.1034-1
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1034-1
Apache Tez	0.9.1.7.1.9.1034-1
Apache Zeppelin	0.8.2.7.1.9.1034-1
Apache ZooKeeper	3.8.1.7.1.9.1034-1

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1034-1
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1034-1
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1034-1
Search	1.0.0.7.1.9.1034-1
Schema Registry	0.10.0.7.1.9.1034-1
Streams Messaging Manager	2.3.0.7.1.9.1034-1
Streams Replication Manager	1.1.0.7.1.9.1034-1

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1034-1
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1034-1
Spark Atlas Connector	0.1.0.7.1.9.1034-1
Spark Schema Registry	1.1.0.7.1.9.1034-1

Cloudera Runtime 7.1.9.1032-3

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 5 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1032-3
Apache Atlas	2.1.0.7.1.9.1032-3
Apache Avro	1.11.1.7.1.9.1032-3
Apache Calcite	1.19.0.7.1.9.1032-3
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1032-3
Apache HBase	2.4.17.7.1.9.1032-3
Apache Hive	3.1.3000.7.1.9.1032-3
Apache Iceberg	1.3.0.7.1.9.1032-3
Apache Impala	4.0.0.7.1.9.1032-3
Apache Kafka	3.4.1.7.1.9.1032-3
Apache Knox	1.3.0.7.1.9.1032-3
Apache Kudu	1.17.0.7.1.9.1032-3
Apache Livy	0.7.2.7.1.9.1032-3
Apache MapReduce	3.1.1.7.1.9.1032-3
Apache Ozone	1.4.0.7.1.9.1032-3
Apache Oozie	5.1.0.7.1.9.1032-3
Apache ORC	1.5.1.7.1.9.1032-3
Apache Parquet	1.10.99.7.1.9.1032-3
Apache Phoenix	5.1.1.7.1.9.1032-3
Apache Ranger	2.4.0.7.1.9.1032-3
Apache Solr	8.11.2.7.1.9.1032-3
Apache Spark 2.x	2.4.8.7.1.9.1032-3
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1032-3
Apache Tez	0.9.1.7.1.9.1032-3
Apache Zeppelin	0.8.2.7.1.9.1032-3
Apache ZooKeeper	3.8.1.7.1.9.1032-3

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1032-3
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1032-3
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1032-3
Search	1.0.0.7.1.9.1032-3
Schema Registry	0.10.0.7.1.9.1032-3
Streams Messaging Manager	2.3.0.7.1.9.1032-3
Streams Replication Manager	1.1.0.7.1.9.1032-3

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1032-3
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1032-3
Spark Atlas Connector	0.1.0.7.1.9.1032-3
Spark Schema Registry	1.1.0.7.1.9.1032-3

Cloudera Runtime 7.1.9.1023-3

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 4 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1023-3
Apache Atlas	2.1.0.7.1.9.1023-3
Apache Avro	1.11.1.7.1.9.1023-3
Apache Calcite	1.19.0.7.1.9.1023-3
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1023-3
Apache HBase	2.4.17.7.1.9.1023-3
Apache Hive	3.1.3000.7.1.9.1023-3
Apache Iceberg	1.3.0.7.1.9.1023-3
Apache Impala	4.0.0.7.1.9.1023-3
Apache Kafka	3.4.1.7.1.9.1023-3
Apache Knox	1.3.0.7.1.9.1023-3
Apache Kudu	1.17.0.7.1.9.1023-3
Apache Livy	0.7.2.7.1.9.1023-3
Apache MapReduce	3.1.1.7.1.9.1023-3
Apache Ozone	1.4.0.7.1.9.1023-3
Apache Oozie	5.1.0.7.1.9.1023-3
Apache ORC	1.5.1.7.1.9.1023-3
Apache Parquet	1.10.99.7.1.9.1023-3
Apache Phoenix	5.1.1.7.1.9.1023-3
Apache Ranger	2.4.0.7.1.9.1023-3
Apache Solr	8.11.2.7.1.9.1023-3
Apache Spark 2.x	2.4.8.7.1.9.1023-3
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1023-3
Apache Tez	0.9.1.7.1.9.1023-3
Apache Zeppelin	0.8.2.7.1.9.1023-3

Component	Version
Apache ZooKeeper	3.8.1.7.1.9.1023-3

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1023-3
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1023-3
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1023-3
Search	1.0.0.7.1.9.1023-3
Schema Registry	0.10.0.7.1.9.1023-3
Streams Messaging Manager	2.3.0.7.1.9.1023-3
Streams Replication Manager	1.1.0.7.1.9.1023-3

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1023-3
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1023-3
Spark Atlas Connector	0.1.0.7.1.9.1023-3
Spark Schema Registry	1.1.0.7.1.9.1023-3

Cloudera Runtime 7.1.9.1015-6

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 3 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1015-6
Apache Atlas	2.1.0.7.1.9.1015-6
Apache Avro	1.11.1.7.1.9.1015-6
Apache Calcite	1.19.0.7.1.9.1015-6
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1015-6
Apache HBase	2.4.17.7.1.9.1015-6
Apache Hive	3.1.3000.7.1.9.1015-6
Apache Iceberg	1.3.0.7.1.9.1015-6
Apache Impala	4.0.0.7.1.9.1015-6
Apache Kafka	3.4.1.7.1.9.1015-6
Apache Knox	1.3.0.7.1.9.1015-6
Apache Kudu	1.17.0.7.1.9.1015-6

Component	Version
Apache Livy	0.7.2.7.1.9.1015-6
Apache MapReduce	3.1.1.7.1.9.1015-6
Apache Ozone	1.4.0.7.1.9.1015-6
Apache Oozie	5.1.0.7.1.9.1015-6
Apache ORC	1.5.1.7.1.9.1015-6
Apache Parquet	1.10.99.7.1.9.1015-6
Apache Phoenix	5.1.1.7.1.9.1015-6
Apache Ranger	2.4.0.7.1.9.1015-6
Apache Solr	8.11.2.7.1.9.1015-6
Apache Spark 2.x	2.4.8.7.1.9.1015-6
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1015-6
Apache Tez	0.9.1.7.1.9.1015-6
Apache Zeppelin	0.8.2.7.1.9.1015-6
Apache ZooKeeper	3.8.1.7.1.9.1015-6

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1015-6
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1015-6
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1015-6
Search	1.0.0.7.1.9.1015-6
Schema Registry	0.10.0.7.1.9.1015-6
Streams Messaging Manager	2.3.0.7.1.9.1015-6
Streams Replication Manager	1.1.0.7.1.9.1015-6

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1015-6
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1015-6
Spark Atlas Connector	0.1.0.7.1.9.1015-6
Spark Schema Registry	1.1.0.7.1.9.1015-6

Cloudera Runtime 7.1.9.1010-2

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 2 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1010-2
Apache Atlas	2.1.0.7.1.9.1010-2
Apache Avro	1.11.1.7.1.9.1010-2
Apache Calcite	1.19.0.7.1.9.1010-2
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1010-2
Apache HBase	2.4.17.7.1.9.1010-2
Apache Hive	3.1.3000.7.1.9.1010-2
Apache Iceberg	1.3.0.7.1.9.1010-2
Apache Impala	4.0.0.7.1.9.1010-2
Apache Kafka	3.4.1.7.1.9.1010-2
Apache Knox	1.3.0.7.1.9.1010-2
Apache Kudu	1.17.0.7.1.9.1010-2
Apache Livy	0.7.2.7.1.9.1010-2
Apache MapReduce	3.1.1.7.1.9.1010-2
Apache Ozone	1.4.0.7.1.9.1010-2
Apache Oozie	5.1.0.7.1.9.1010-2
Apache ORC	1.5.1.7.1.9.1010-2
Apache Parquet	1.10.99.7.1.9.1010-2
Apache Phoenix	5.1.1.7.1.9.1010-2
Apache Ranger	2.4.0.7.1.9.1010-2
Apache Solr	8.11.2.7.1.9.1010-2
Apache Spark 2.x	2.4.8.7.1.9.1010-2
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1010-2
Apache Tez	0.9.1.7.1.9.1010-2
Apache Zeppelin	0.8.2.7.1.9.1010-2
Apache ZooKeeper	3.8.1.7.1.9.1010-2

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1010-2
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1010-2
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1010-2
Search	1.0.0.7.1.9.1010-2
Schema Registry	0.10.0.7.1.9.1010-2
Streams Messaging Manager	2.3.0.7.1.9.1010-2
Streams Replication Manager	1.1.0.7.1.9.1010-2

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1010-2
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1010-2
Spark Atlas Connector	0.1.0.7.1.9.1010-2
Spark Schema Registry	1.1.0.7.1.9.1010-2

Cloudera Runtime 7.1.9.1005-2

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 CHF 1 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1005-2
Apache Atlas	2.1.0.7.1.9.1005-2
Apache Avro	1.11.1.7.1.9.1005-2
Apache Calcite	1.19.0.7.1.9.1005-2
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1005-2
Apache HBase	2.4.17.7.1.9.1005-2
Apache Hive	3.1.3000.7.1.9.1005-2
Apache Iceberg	1.3.0.7.1.9.1005-2
Apache Impala	4.0.0.7.1.9.1005-2
Apache Kafka	3.4.1.7.1.9.1005-2
Apache Knox	1.3.0.7.1.9.1005-2
Apache Kudu	1.17.0.7.1.9.1005-2
Apache Livy	0.7.2.7.1.9.1005-2
Apache MapReduce	3.1.1.7.1.9.1005-2
Apache Ozone	1.4.0.7.1.9.1005-2
Apache Oozie	5.1.0.7.1.9.1005-2
Apache ORC	1.5.1.7.1.9.1005-2
Apache Parquet	1.10.99.7.1.9.1005-2
Apache Phoenix	5.1.1.7.1.9.1005-2
Apache Ranger	2.4.0.7.1.9.1005-2
Apache Solr	8.11.2.7.1.9.1005-2
Apache Spark 2.x	2.4.8.7.1.9.1005-2
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1005-2
Apache Tez	0.9.1.7.1.9.1005-2
Apache Zeppelin	0.8.2.7.1.9.1005-2

Component	Version
Apache ZooKeeper	3.8.1.7.1.9.1005-2

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1005-2
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1005-2
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1005-2
Search	1.0.0.7.1.9.1005-2
Schema Registry	0.10.0.7.1.9.1005-2
Streams Messaging Manager	2.3.0.7.1.9.1005-2
Streams Replication Manager	1.1.0.7.1.9.1005-2

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1005-2
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1005-2
Spark Atlas Connector	0.1.0.7.1.9.1005-2
Spark Schema Registry	1.1.0.7.1.9.1005-2

Cloudera Runtime 7.1.9.1000-103

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 SP1 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.1.9.1000-103
Apache Atlas	2.1.0.7.1.9.1000-103
Apache Avro	1.11.1.7.1.9.1000-103
Apache Calcite	1.19.0.7.1.9.1000-103
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.1.9.1000-103
Apache HBase	2.4.17.7.1.9.1000-103
Apache Hive	3.1.3000.7.1.9.1000-103
Apache Iceberg	1.3.0.7.1.9.1000-103
Apache Impala	4.0.0.7.1.9.1000-103
Apache Kafka	3.4.1.7.1.9.1000-103
Apache Knox	1.3.0.7.1.9.1000-103
Apache Kudu	1.17.0.7.1.9.1000-103

Component	Version
Apache Livy	0.7.2.7.1.9.1000-103
Apache MapReduce	3.1.1.7.1.9.1000-103
Apache Ozone	1.4.0.7.1.9.1000-103
Apache Oozie	5.1.0.7.1.9.1000-103
Apache ORC	1.5.1.7.1.9.1000-103
Apache Parquet	1.10.99.7.1.9.1000-103
Apache Phoenix	5.1.1.7.1.9.1000-103
Apache Ranger	2.4.0.7.1.9.1000-103
Apache Solr	8.11.2.7.1.9.1000-103
Apache Spark 2.x	2.4.8.7.1.9.1000-103
Apache Spark 3.x	3.3.7191000.3-1 CDS
Apache Sqoop	1.4.7.7.1.9.1000-103
Apache Tez	0.9.1.7.1.9.1000-103
Apache Zeppelin	0.8.2.7.1.9.1000-103
Apache ZooKeeper	3.8.1.7.1.9.1000-103

Other Components

Component	Version
Cruise Control	2.5.116.7.1.9.1000-103
Data Analytics Studio	1.4.2
GCS Connector	2.1.2.7.1.9.1000-103
HBase Indexer	1.5.0.7.5.0-141
Hue	4.11.0.7.1.9.1000-103
Search	1.0.0.7.1.9.1000-103
Schema Registry	0.10.0.7.1.9.1000-103
Streams Messaging Manager	2.3.0.7.1.9.1000-103
Streams Replication Manager	1.1.0.7.1.9.1000-103

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.1.9.1000-103
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0.7.1.9.1000-103
Spark Atlas Connector	0.1.0.7.1.9.1000-103
Spark Schema Registry	1.1.0.7.1.9.1000-103

Using the Cloudera Runtime Maven repository 7.1.9 SP1

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 repository.cloudera.com.



Important: When you build an application JAR, do not include CDH JARs, because they are already provided. If you do, upgrading CDH can break your application. To avoid this situation, set the Maven dependency scope to provided. If you have already built applications which include the CDH JARs, update the dependency to set scope to provided and recompile.

The following is a sample POM (pom.xml) file:

```
<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/maven-v4_0_0.xsd">
  <repositories>
    <repository>
      <id>cloudera</id>
      <url>https://repository.cloudera.com/artifactory/cloudera-repos/</url>
    </repository>
  </repositories>
</project>
```

Runtime 7.1.9.1073-8

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	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-classification-updater	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-client-common	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-client-v1	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-client-v2	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-common	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-distro	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-docs	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-graphdb-api	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-graphdb-common	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-intg	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-notification	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-repository	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-server-api	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	atlas-testtools	2.1.0.7.1.9.1073-8

Project	groupId	artifactId	version
Atlas	org.apache.atlas	hbase-bridge	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	hbase-bridge-shim	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	hbase-testing-util	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	hdfs-model	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	hive-bridge	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	hive-bridge-shim	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	impala-bridge	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	impala-bridge-shim	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	impala-hook-api	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	kafka-bridge	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	kafka-bridge-shim	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	navigator-to-atlas	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	sample-app	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	sqoop-bridge	2.1.0.7.1.9.1073-8
Atlas	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.1.9.1073-8
Avro	org.apache.avro	avro	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1073-8
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1073-8
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1073-8
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1073-8
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1073-8
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1073-8
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1073-8

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1073-8

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1073-8
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1073-8
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1073-8

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1073-8
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1073-8
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1073-8

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1073-8
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1073-8
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1073-8
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1073-8

Project	groupId	artifactId	version
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1073-8

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1073-8
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1073-8
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1073-8

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-knoxssso	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-knoxssout	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1073-8

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1073-8
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1073-8
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1073-8
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1073-8
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1073-8
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1073-8
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1073-8
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1073-8
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1073-8
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1073-8
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1073-8
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1073-8
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1073-8
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1073-8
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1073-8
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1073-8
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1073-8
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1073-8
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1073-8

Project	groupId	artifactId	version
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1073-8
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1073-8
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1073-8
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1073-8
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1073-8
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-analyzers-openslp	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1073-8
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1073-8

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1073-8
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1073-8
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1073-8
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1073-8
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1073-8
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1073-8
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-managed-rocksdbs	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1073-8

Project	groupId	artifactId	version
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-filesystem	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-filesystem-common	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-filesystem-hadoop2	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-filesystem-hadoop3	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-filesystem-shaded	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-httpfgateway	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1073-8
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1073-8

Project	groupId	artifactId	version
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1073-8
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compatible	6.0.0.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.1.6	5.1.1.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.2.5	5.1.1.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.3.0	5.1.1.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.0	5.1.1.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.1	5.1.1.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1073-8
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1073-8
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1073-8

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1073-8

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-trino-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1073-8
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1073-8
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1073-8
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1073-8
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1073-8

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1073-8
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1073-8
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1073-8
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1073-8
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1073-8

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1073-8
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1073-8
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1073-8
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1073-8
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1073-8
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1073-8
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1073-8
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1073-8
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1073-8
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1073-8
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1073-8
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1073-8
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1073-8
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1073-8
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1073-8
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1073-8
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.1.9.1073-8
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1073-8
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1073-8
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1073-8
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1073-8
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1073-8
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1073-8

Runtime 7.1.9.1069-3

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	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-classification-updater	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-client-common	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-client-v1	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-client-v2	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-common	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-distro	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-docs	2.1.0.7.1.9.1069-3

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-graphdb-api	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-graphdb-common	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-intg	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-notification	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-repository	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-server-api	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	atlas-testtools	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	hbase-bridge	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	hbase-bridge-shim	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	hbase-testing-util	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	hdfs-model	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	hive-bridge	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	hive-bridge-shim	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	impala-bridge	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	impala-bridge-shim	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	impala-hook-api	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	kafka-bridge	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	kafka-bridge-shim	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	navigator-to-atlas	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	sample-app	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	sqoop-bridge	2.1.0.7.1.9.1069-3
Atlas	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.1.9.1069-3
Avro	org.apache.avro	avro	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1069-3

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1069-3
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1069-3
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1069-3
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1069-3
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1069-3
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1069-3
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1069-3

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-miniclient	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1069-3

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1069-3
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1069-3
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1069-3

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1069-3
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1069-3
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1069-3

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1069-3
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1069-3
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1069-3
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1069-3

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1069-3
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1069-3
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1069-3

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-knoxsso	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-knoxssout	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1069-3

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1069-3
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1069-3
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1069-3
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1069-3
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1069-3
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1069-3
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1069-3

Project	groupId	artifactId	version
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1069-3
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1069-3
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1069-3
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1069-3
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1069-3
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1069-3
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1069-3
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1069-3
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1069-3
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1069-3
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1069-3
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1069-3
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1069-3
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1069-3
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1069-3
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1069-3
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1069-3
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-analyzers-opennlp	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1069-3

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1069-3
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1069-3
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1069-3
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1069-3
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1069-3
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1069-3
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1069-3
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1069-3

Project	groupId	artifactId	version
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-filessystem	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-filessystem-common	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-filessystem-hadoop2	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-filessystem-hadoop3	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-filessystem-shaded	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-httpsgateway	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1069-3
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1069-3

Project	groupId	artifactId	version
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1069-3
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compatible	6.0.0.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.1.6	5.1.1.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.2.5	5.1.1.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.3.0	5.1.1.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.0	5.1.1.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.1	5.1.1.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1069-3
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1069-3
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1069-3

Project	groupId	artifactId	version
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1069-3

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-trino-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1069-3
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1069-3
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1069-3
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1069-3
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1069-3
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1069-3
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1069-3
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1069-3
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1069-3
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1069-3

Project	groupId	artifactId	version
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1069-3
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1069-3
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1069-3
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1069-3
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1069-3
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1069-3
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1069-3
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1069-3
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1069-3
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1069-3
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1069-3
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1069-3
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1069-3

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1069-3
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1069-3
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1069-3
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1069-3
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1069-3
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1069-3
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1069-3
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1069-3
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1069-3
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1069-3
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1069-3
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1069-3
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1069-3
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1069-3
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1069-3
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1069-3
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.1.9.1069-3
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1069-3
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1069-3
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1069-3
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1069-3
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1069-3
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1069-3

Runtime 7.1.9.1066-8

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	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-classification-updater	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-client-common	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-client-v1	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-client-v2	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-common	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-distro	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-docs	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-graphdb-api	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-graphdb-common	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-intg	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-notification	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-repository	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-server-api	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	atlas-testtools	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	hbase-bridge	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	hbase-bridge-shim	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	hbase-testing-util	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	hdfs-model	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	hive-bridge	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	hive-bridge-shim	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	impala-bridge	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	impala-bridge-shim	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	impala-hook-api	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	kafka-bridge	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	kafka-bridge-shim	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	navigator-to-atlas	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	sample-app	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	sqoop-bridge	2.1.0.7.1.9.1066-8
Atlas	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.1.9.1066-8
Avro	org.apache.avro	avro	1.11.1.7.1.9.1066-8

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1066-8
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1066-8
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1066-8
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1066-8
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1066-8
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1066-8
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1066-8

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1066-8

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1066-8
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1066-8
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1066-8

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1066-8
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1066-8
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1066-8

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1066-8
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1066-8
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1066-8
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1066-8

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1066-8
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1066-8
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1066-8

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1066-8

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-knoxsso	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-knoxsout	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1066-8

Project	groupId	artifactId	version
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1066-8
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1066-8
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1066-8
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1066-8
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1066-8
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1066-8
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1066-8
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1066-8
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1066-8
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1066-8
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1066-8
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1066-8
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1066-8
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1066-8
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1066-8
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1066-8
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1066-8
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1066-8
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1066-8
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1066-8
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1066-8
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1066-8
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1066-8
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1066-8
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-analyzers-openssl	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1066-8

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1066-8
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1066-8
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1066-8

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1066-8
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1066-8
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1066-8
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1066-8
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1066-8
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-filesystem	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-filesystem-common	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-filesystem-hadoop2	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-filesystem-hadoop3	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-filesystem-shaded	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-httpsgateway	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1066-8

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1066-8
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1066-8
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compat	6.0.0.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.1.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.1.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.1.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.1.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.1.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1066-8

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1066-8
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1066-8
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1066-8

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-trino-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1066-8

Project	groupId	artifactId	version
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1066-8
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1066-8
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1066-8
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1066-8
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1066-8

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1066-8
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1066-8
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1066-8
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1066-8
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1066-8
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1066-8
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1066-8
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1066-8
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1066-8
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1066-8
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1066-8
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1066-8
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1066-8
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1066-8
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1066-8
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1066-8
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1066-8
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1066-8

Project	groupId	artifactId	version
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1066-8
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1066-8
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.1.9.1066-8
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1066-8
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1066-8
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1066-8
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1066-8
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1066-8
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1066-8

Runtime 7.1.9.1064-1

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	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-classification-updater	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-client-common	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-client-v1	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-client-v2	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-common	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-distro	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-docs	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-graphdb-api	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-graphdb-common	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-intg	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-notification	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-repository	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-server-api	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	atlas-testtools	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	hbase-bridge	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	hbase-bridge-shim	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	hbase-testing-util	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	hdfs-model	2.1.0.7.1.9.1064-1

Project	groupId	artifactId	version
Atlas	org.apache.atlas	hive-bridge	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	hive-bridge-shim	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	impala-bridge	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	impala-bridge-shim	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	impala-hook-api	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	kafka-bridge	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	kafka-bridge-shim	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	navigator-to-atlas	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	sample-app	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	sqoop-bridge	2.1.0.7.1.9.1064-1
Atlas	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.1.9.1064-1
Avro	org.apache.avro	avro	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1064-1
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1064-1
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1064-1
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1064-1
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1064-1
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1064-1
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1064-1

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1064-1

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1064-1
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1064-1
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1064-1

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1064-1
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1064-1
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1064-1

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1064-1
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1064-1
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1064-1
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1064-1

Project	groupId	artifactId	version
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1064-1

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1064-1
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1064-1
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1064-1

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-knoxsso	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-knoxssout	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1064-1

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1064-1
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1064-1
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1064-1
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1064-1
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1064-1
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1064-1
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1064-1
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1064-1
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1064-1
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1064-1
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1064-1
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1064-1
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1064-1
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1064-1
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1064-1
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1064-1
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1064-1
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1064-1
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1064-1
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1064-1
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1064-1
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1064-1
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1064-1

Project	groupId	artifactId	version
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1064-1
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-analyzers-opennlp	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1064-1
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1064-1

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1064-1
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1064-1
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1064-1
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1064-1
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1064-1
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1064-1
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1064-1

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-filessystem	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-filessystem-common	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-filessystem-hadoop2	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-filessystem-hadoop3	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-filessystem-shaded	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-httpfsgateway	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1064-1
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1064-1
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1064-1

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compat	6.0.0.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.1.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.1.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.1.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.1.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.1.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1064-1
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1064-1
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1064-1

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1064-1

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-trino-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1064-1
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1064-1
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1064-1
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1064-1
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1064-1

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1064-1
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1064-1
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1064-1
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1064-1
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1064-1
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1064-1

Project	groupId	artifactId	version
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1064-1
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1064-1
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1064-1
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1064-1
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1064-1
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1064-1
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1064-1
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1064-1
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1064-1
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1064-1
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1064-1
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1064-1
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1064-1
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1064-1
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinpector	3.8.1.7.1.9.1064-1
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1064-1
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1064-1
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1064-1
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1064-1
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1064-1
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1064-1

Runtime 7.1.9.1059-4

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	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-classification-updater	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-client-common	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-client-v1	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-client-v2	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-common	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-distro	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-docs	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-graphdb-api	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-graphdb-common	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.1.0.7.1.9.1059-4

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-intg	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-notification	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-repository	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-server-api	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	atlas-testtools	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	hbase-bridge	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	hbase-bridge-shim	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	hbase-testing-util	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	hdfs-model	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	hive-bridge	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	hive-bridge-shim	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	impala-bridge	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	impala-bridge-shim	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	impala-hook-api	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	kafka-bridge	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	kafka-bridge-shim	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	navigator-to-atlas	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	sample-app	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	sqoop-bridge	2.1.0.7.1.9.1059-4
Atlas	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.1.9.1059-4
Avro	org.apache.avro	avro	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1059-4

Project	groupId	artifactId	version
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1059-4
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1059-4
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1059-4
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1059-4
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1059-4
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1059-4
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1059-4

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-miniclusterc	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1059-4
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1059-4

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1059-4
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1059-4
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1059-4
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1059-4

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1059-4
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1059-4

Project	groupId	artifactId	version
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1059-4
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1059-4

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1059-4
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1059-4
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1059-4

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-knoxsso	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-knoxssout	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1059-4

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1059-4
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1059-4
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1059-4
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1059-4
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1059-4
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1059-4
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1059-4
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1059-4
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1059-4
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1059-4
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1059-4

Project	groupId	artifactId	version
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1059-4
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1059-4
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1059-4
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1059-4
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1059-4
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1059-4
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1059-4
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1059-4
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1059-4
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1059-4
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1059-4
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1059-4
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1059-4
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-analyzers-openslp	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1059-4

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1059-4
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1059-4
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1059-4
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1059-4
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1059-4
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1059-4
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1059-4
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1059-4

Project	groupId	artifactId	version
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-filesystem	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-filesystem-common	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-filesystem-hadoop2	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-filesystem-hadoop3	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-filesystem-shaded	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-httpsgateway	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1059-4
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1059-4

Project	groupId	artifactId	version
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1059-4
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compatible	6.0.0.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.1.6	5.1.1.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.2.5	5.1.1.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.3.0	5.1.1.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.0	5.1.1.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.1	5.1.1.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1059-4
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1059-4
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1059-4

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1059-4

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-trino-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1059-4
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1059-4
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1059-4
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1059-4
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1059-4
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1059-4
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1059-4
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1059-4
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1059-4
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1059-4
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1059-4
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1059-4
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1059-4
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1059-4

Project	groupId	artifactId	version
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1059-4
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1059-4
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1059-4
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1059-4
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1059-4
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1059-4
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1059-4
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1059-4
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1059-4

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1059-4
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1059-4
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1059-4
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1059-4
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1059-4
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1059-4
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1059-4
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1059-4
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1059-4
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1059-4
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1059-4
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1059-4
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1059-4
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1059-4
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1059-4
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1059-4
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.1.9.1059-4
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1059-4
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1059-4
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1059-4
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1059-4
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1059-4
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1059-4

Runtime 7.1.9.1054-4

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	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.1.9.1054-4

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-classification-updater	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-client-common	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-client-v1	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-client-v2	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-common	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-distro	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-intg	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-notification	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-repository	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-server-api	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	atlas-testtools	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	hbase-bridge	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	hbase-bridge-shim	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	hbase-testing-util	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	hdfs-model	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	hive-bridge	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	hive-bridge-shim	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	impala-bridge	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	impala-bridge-shim	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	impala-hook-api	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	kafka-bridge	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	kafka-bridge-shim	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	navigator-to-atlas	2.1.0.7.1.9.1054-4
Atlas	org.apache.atlas	sample-app	2.1.0.7.1.9.1054-4
Avro	org.apache.avro	avro	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1054-4

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1054-4
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1054-4
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1054-4
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1054-4
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1054-4
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1054-4
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1054-4

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1054-4

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1054-4
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1054-4
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1054-4

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1054-4
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1054-4
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1054-4

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1054-4
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1054-4
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1054-4
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1054-4

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1054-4
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1054-4
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1054-4

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-knoxssso	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-knoxssout	1.3.0.7.1.9.1054-4

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1054-4
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1054-4
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1054-4
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1054-4
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1054-4
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1054-4

Project	groupId	artifactId	version
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1054-4
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1054-4
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1054-4
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1054-4
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1054-4
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1054-4
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1054-4
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1054-4
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1054-4
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1054-4
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1054-4
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1054-4
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1054-4
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1054-4
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1054-4
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1054-4
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1054-4
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1054-4
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-analyzers-openslp	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1054-4

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1054-4
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1054-4
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1054-4
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1054-4
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1054-4
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1054-4
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1054-4
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1054-4

Project	groupId	artifactId	version
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-managed-rocksdbs	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-filesystem	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-filesystem-common	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-filesystem-hadoop2	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-filesystem-hadoop3	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-filesystem-shaded	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-httpfgateway	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1054-4
Ozone	org.apache.ozone	rocksdbs-checkpoint-differ	1.4.0.7.1.9.1054-4

Project	groupId	artifactId	version
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1054-4
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compatible	6.0.0.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.1.6	5.1.1.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.2.5	5.1.1.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.3.0	5.1.1.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.0	5.1.1.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.1	5.1.1.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1054-4
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1054-4
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1054-4

Project	groupId	artifactId	version
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1054-4

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-trino-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1054-4
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1054-4
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1054-4
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1054-4
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1054-4
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1054-4
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1054-4
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1054-4
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1054-4

Project	groupId	artifactId	version
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1054-4
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1054-4
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1054-4
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1054-4
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1054-4
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1054-4
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1054-4
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1054-4
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1054-4
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1054-4
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1054-4
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1054-4
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1054-4
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1054-4

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1054-4
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1054-4
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1054-4
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1054-4
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1054-4
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1054-4
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1054-4
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1054-4
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1054-4
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1054-4
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1054-4
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1054-4
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1054-4
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1054-4
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1054-4
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1054-4
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.1.9.1054-4
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1054-4
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1054-4
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1054-4
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1054-4
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1054-4
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1054-4

Runtime 7.1.9.1045-5

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	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-classification-updater	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-client-common	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-client-v1	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-client-v2	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-common	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-distro	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-docs	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-graphdb-api	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-graphdb-common	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-intg	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-notification	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-repository	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-server-api	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	atlas-testtools	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	hbase-bridge	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	hbase-bridge-shim	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	hbase-testing-util	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	hdfs-model	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	hive-bridge	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	hive-bridge-shim	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	impala-bridge	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	impala-bridge-shim	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	impala-hook-api	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	kafka-bridge	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	kafka-bridge-shim	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	navigator-to-atlas	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	sample-app	2.1.0.7.1.9.1045-5
Atlas	org.apache.atlas	sqoop-bridge	2.1.0.7.1.9.1045-5

Project	groupId	artifactId	version
Atlas	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.1.9.1045-5
Avro	org.apache.avro	avro	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1045-5
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1045-5
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1045-5
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1045-5
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1045-5
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1045-5
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1045-5

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1045-5

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1045-5
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1045-5
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1045-5

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1045-5
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1045-5
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1045-5

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1045-5
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1045-5
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1045-5
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1045-5

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1045-5
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1045-5
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1045-5

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1045-5

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-knoxsso	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-knoxsout	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1045-5

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1045-5
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1045-5
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1045-5
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1045-5
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1045-5
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1045-5
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1045-5
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1045-5
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1045-5
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1045-5
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1045-5
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1045-5
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1045-5
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1045-5
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1045-5
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1045-5
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1045-5
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1045-5
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1045-5
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1045-5
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1045-5
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1045-5
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1045-5
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1045-5
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-analyzers-opennlp	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-analyzers-smartcnn	8.11.2.7.1.9.1045-5

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1045-5
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1045-5

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1045-5
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1045-5
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1045-5
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1045-5
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1045-5
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1045-5
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-filesystem	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-filesystem-common	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-filesystem-hadoop2	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-filesystem-hadoop3	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-filesystem-shaded	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-httpsgateway	1.4.0.7.1.9.1045-5

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1045-5
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1045-5
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compat	6.0.0.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.1.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.1.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.1.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.1.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.1.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1045-5

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1045-5
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1045-5
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1045-5

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-trino-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1045-5

Project	groupId	artifactId	version
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1045-5
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1045-5
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1045-5
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1045-5
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1045-5

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1045-5
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1045-5
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1045-5
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1045-5
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1045-5
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1045-5
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1045-5
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1045-5
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1045-5
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1045-5
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1045-5
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1045-5
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1045-5
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1045-5
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1045-5
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1045-5

Project	groupId	artifactId	version
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1045-5
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1045-5
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1045-5
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1045-5
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.1.9.1045-5
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1045-5
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1045-5
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1045-5
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1045-5
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1045-5
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1045-5

Runtime 7.1.9.1042-1

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	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-classification-updater	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-client-common	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-client-v1	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-client-v2	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-common	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-distro	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-docs	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-graphdb-api	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-graphdb-common	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-intg	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-notification	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-repository	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-server-api	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	atlas-testtools	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	hbase-bridge	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	hbase-bridge-shim	2.1.0.7.1.9.1042-1

Project	groupId	artifactId	version
Atlas	org.apache.atlas	hbase-testing-util	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	hdfs-model	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	hive-bridge	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	hive-bridge-shim	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	impala-bridge	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	impala-bridge-shim	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	impala-hook-api	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	kafka-bridge	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	kafka-bridge-shim	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	navigator-to-atlas	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	sample-app	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	sqoop-bridge	2.1.0.7.1.9.1042-1
Atlas	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.1.9.1042-1
Avro	org.apache.avro	avro	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1042-1
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1042-1
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1042-1
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1042-1
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1042-1
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1042-1
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1042-1

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.1.9.1042-1

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1042-1
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1042-1
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1042-1

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1042-1
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1042-1
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1042-1

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1042-1
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1042-1
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1042-1
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1042-1

Project	groupId	artifactId	version
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1042-1

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1042-1
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1042-1
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1042-1

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-knoxsso	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-knoxsout	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1042-1

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1042-1
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1042-1
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1042-1
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1042-1
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1042-1
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1042-1
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1042-1
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1042-1
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1042-1
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1042-1
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1042-1
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1042-1
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1042-1
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1042-1
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1042-1
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1042-1
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1042-1
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1042-1
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1042-1
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1042-1
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1042-1

Project	groupId	artifactId	version
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1042-1
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1042-1
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1042-1
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-analyzers-openslp	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1042-1
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1042-1

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1042-1
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1042-1
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1042-1
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1042-1
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1042-1
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1042-1
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1042-1

Project	groupId	artifactId	version
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-filesystem	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-filesystem-common	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-filesystem-hadoop2	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-filesystem-hadoop3	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-filesystem-shaded	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-httpsgateway	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1042-1
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1042-1

Project	groupId	artifactId	version
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1042-1
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compatible	6.0.0.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.1.6	5.1.1.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.2.5	5.1.1.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.3.0	5.1.1.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.0	5.1.1.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.1	5.1.1.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1042-1
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1042-1
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1042-1

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1042-1

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1042-1
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1042-1
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1042-1
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1042-1
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1042-1

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1042-1
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1042-1
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1042-1
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1042-1
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1042-1
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1042-1

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1042-1
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1042-1
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1042-1
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1042-1
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1042-1
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1042-1
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1042-1
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1042-1
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1042-1
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1042-1
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1042-1
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1042-1
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1042-1
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1042-1
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1042-1
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.1.9.1042-1
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1042-1
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1042-1
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1042-1
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1042-1
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1042-1
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1042-1

Runtime 7.1.9.1037-2

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	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-aws-s3-bridge	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-azure-adls-bridge	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-classification-updater	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-client-common	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-client-v1	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-client-v2	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-client-v2-shaded	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-common	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-distro	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-docs	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-graphdb-api	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-graphdb-common	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-graphdb-janus	3.0.0.7.1.9.1037-2

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-hdfs-bridge	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-index-repair-tool	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-intg	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-notification	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-plugin-classloader	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-repository	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-server-api	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	atlas-testtools	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	hbase-bridge	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	hbase-bridge-shim	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	hbase-testing-util	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	hdfs-model	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	hive-bridge	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	hive-bridge-shim	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	impala-bridge	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	impala-bridge-shim	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	impala-hook-api	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	kafka-bridge	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	kafka-bridge-shim	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	navigator-to-atlas	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	sample-app	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	sqoop-bridge	3.0.0.7.1.9.1037-2
Atlas	org.apache.atlas	sqoop-bridge-shim	3.0.0.7.1.9.1037-2
Avro	org.apache.avro	avro	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1037-2

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1037-2
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1037-2
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1037-2
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1037-2
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1037-2
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1037-2
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1037-2

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1037-2

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1037-2
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1037-2
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1037-2
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1037-2
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1037-2

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1037-2
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1037-2

Project	groupId	artifactId	version
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1037-2
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1037-2
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1037-2

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1037-2
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1037-2
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1037-2

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-knoxssso	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-knoxssout	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1037-2

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1037-2
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1037-2
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1037-2
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1037-2
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1037-2
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1037-2
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1037-2
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1037-2
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1037-2
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1037-2

Project	groupId	artifactId	version
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1037-2
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1037-2
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1037-2
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1037-2
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1037-2
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1037-2
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1037-2
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1037-2
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1037-2
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1037-2
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1037-2
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1037-2
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1037-2
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1037-2
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-analyzers-opennlp	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1037-2

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1037-2
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1037-2
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1037-2
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1037-2
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1037-2
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1037-2
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1037-2
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1037-2

Project	groupId	artifactId	version
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-filessystem	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-filessystem-common	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-filessystem-hadoop2	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-filessystem-hadoop3	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-filessystem-shaded	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-httpsgateway	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1037-2
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1037-2

Project	groupId	artifactId	version
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1037-2
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compatible	6.0.0.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.1.6	5.1.1.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.2.5	5.1.1.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.3.0	5.1.1.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.0	5.1.1.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.1	5.1.1.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1037-2
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1037-2
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1037-2

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1037-2

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1037-2
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1037-2
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1037-2
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1037-2
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1037-2
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1037-2
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1037-2
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1037-2
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1037-2
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1037-2
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1037-2
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1037-2
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1037-2
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1037-2

Project	groupId	artifactId	version
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1037-2
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1037-2
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1037-2
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1037-2
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1037-2
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1037-2
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1037-2
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1037-2
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1037-2

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1037-2
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1037-2
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1037-2
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1037-2
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1037-2
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1037-2
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1037-2
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1037-2
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1037-2
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1037-2
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1037-2
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1037-2
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1037-2
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1037-2
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1037-2
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1037-2
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.1.9.1037-2
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1037-2
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1037-2
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1037-2
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1037-2
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1037-2
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1037-2

Runtime 7.1.9.1034-1

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	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-aws-s3-bridge	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-azure-adls-bridge	3.0.0.7.1.9.1034-1

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-classification-updater	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-client-common	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-client-v1	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-client-v2	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-client-v2-shaded	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-common	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-distro	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-docs	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-graphdb-api	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-graphdb-common	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-graphdb-janus	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-hdfs-bridge	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-index-repair-tool	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-intg	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-notification	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-plugin-classloader	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-repository	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-server-api	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	atlas-testtools	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	hbase-bridge	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	hbase-bridge-shim	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	hbase-testing-util	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	hdfs-model	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	hive-bridge	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	hive-bridge-shim	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	impala-bridge	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	impala-bridge-shim	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	impala-hook-api	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	kafka-bridge	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	kafka-bridge-shim	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	navigator-to-atlas	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	sample-app	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	sqoop-bridge	3.0.0.7.1.9.1034-1
Atlas	org.apache.atlas	sqoop-bridge-shim	3.0.0.7.1.9.1034-1
Avro	org.apache.avro	avro	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1034-1

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1034-1
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1034-1
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1034-1
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1034-1
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1034-1
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1034-1
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1034-1

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1034-1

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1034-1
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1034-1
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1034-1

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1034-1
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1034-1
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1034-1

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1034-1
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1034-1
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1034-1
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1034-1

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1034-1
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1034-1
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1034-1

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1034-1

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-knoxsso	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-knoxsout	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1034-1

Project	groupId	artifactId	version
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1034-1
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1034-1
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1034-1
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1034-1
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1034-1
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1034-1
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1034-1
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1034-1
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1034-1
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1034-1
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1034-1
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1034-1
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1034-1
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1034-1
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1034-1
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1034-1
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1034-1
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1034-1
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1034-1
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1034-1
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1034-1
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1034-1
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1034-1
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1034-1
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-analyzers-openslp	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1034-1

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1034-1
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1034-1
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1034-1
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1034-1
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1034-1

Project	groupId	artifactId	version
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1034-1
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1034-1
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-filessystem	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-filessystem-common	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-filessystem-hadoop2	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-filessystem-hadoop3	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-filessystem-shaded	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-httpsgateway	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1034-1

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1034-1
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1034-1
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compat	6.0.0.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.1.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.1.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.1.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.1.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.1.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1034-1

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1034-1
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1034-1
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1034-1

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1034-1
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1034-1
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1034-1

Project	groupId	artifactId	version
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1034-1
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1034-1
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1034-1
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1034-1
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1034-1
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1034-1
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1034-1
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1034-1
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1034-1
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1034-1
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1034-1
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1034-1
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1034-1
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1034-1
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1034-1
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1034-1
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1034-1

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1034-1
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1034-1
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1034-1
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1034-1
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1034-1
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1034-1
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1034-1
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1034-1
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1034-1
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1034-1
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1034-1
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1034-1
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1034-1
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1034-1
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1034-1
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1034-1
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1034-1
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1034-1
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1034-1
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.1.9.1034-1
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1034-1

Project	groupId	artifactId	version
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1034-1
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1034-1
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1034-1
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1034-1
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1034-1

Runtime 7.1.9.1032-3

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	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-aws-s3-bridge	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-azure-adls-bridge	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-classification-updater	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-client-common	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-client-v1	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-client-v2	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-client-v2-shaded	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-common	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-distro	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-docs	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-graphdb-api	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-graphdb-common	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-graphdb-janus	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-hdfs-bridge	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-index-repair-tool	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-intg	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-notification	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-plugin-classloader	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-repository	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-server-api	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	atlas-testtools	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	hbase-bridge	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	hbase-bridge-shim	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	hbase-testing-util	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	hdfs-model	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	hive-bridge	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	hive-bridge-shim	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	impala-bridge	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	impala-bridge-shim	3.0.0.7.1.9.1032-3

Project	groupId	artifactId	version
Atlas	org.apache.atlas	impala-hook-api	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	kafka-bridge	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	kafka-bridge-shim	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	navigator-to-atlas	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	sample-app	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	sqoop-bridge	3.0.0.7.1.9.1032-3
Atlas	org.apache.atlas	sqoop-bridge-shim	3.0.0.7.1.9.1032-3
Avro	org.apache.avro	avro	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1032-3
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1032-3
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1032-3
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1032-3
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1032-3
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1032-3
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1032-3

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1032-3

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1032-3
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1032-3
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1032-3

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1032-3
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1032-3
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1032-3

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1032-3
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1032-3
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1032-3
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1032-3

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1032-3

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1032-3
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1032-3
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1032-3

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-knoxsso	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-knoxssout	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1032-3

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1032-3
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1032-3
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1032-3
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1032-3
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1032-3
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1032-3
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1032-3
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1032-3
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1032-3
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1032-3
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1032-3
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1032-3
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1032-3
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1032-3
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1032-3
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1032-3
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1032-3
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1032-3
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1032-3
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1032-3
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1032-3
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1032-3
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1032-3
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1032-3
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1032-3

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-analyzers-openslp	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1032-3
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1032-3

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1032-3
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1032-3
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1032-3
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1032-3
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1032-3
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1032-3
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1032-3

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-filesystem	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-filesystem-common	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-filesystem-hadoop2	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-filesystem-hadoop3	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-filesystem-shaded	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-httpfsgateway	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1032-3
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1032-3
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compatible	6.0.0.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1032-3

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.1.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.1.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.1.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.1.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.1.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1032-3
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1032-3
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1032-3

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1032-3

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1032-3
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1032-3
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1032-3
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1032-3
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1032-3

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1032-3
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1032-3
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1032-3
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1032-3
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1032-3
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1032-3
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1032-3
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1032-3
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1032-3
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1032-3
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1032-3

Project	groupId	artifactId	version
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1032-3
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1032-3
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1032-3
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1032-3
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1032-3
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1032-3
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1032-3
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1032-3
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1032-3
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.1.9.1032-3
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1032-3
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1032-3
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1032-3
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1032-3
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1032-3
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1032-3

Runtime 7.1.9.1023-3

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	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-aws-s3-bridge	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-azure-adls-bridge	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-classification-updater	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-client-common	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-client-v1	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-client-v2	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-client-v2-shaded	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-common	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-distro	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-docs	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-graphdb-api	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-graphdb-common	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-graphdb-janus	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-hdfs-bridge	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-index-repair-tool	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-intg	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-notification	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-plugin-classloader	3.0.0.7.1.9.1023-3

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-repository	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-server-api	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	atlas-testtools	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	hbase-bridge	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	hbase-bridge-shim	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	hbase-testing-util	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	hdfs-model	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	hive-bridge	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	hive-bridge-shim	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	impala-bridge	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	impala-bridge-shim	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	impala-hook-api	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	kafka-bridge	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	kafka-bridge-shim	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	navigator-to-atlas	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	sample-app	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	sqoop-bridge	3.0.0.7.1.9.1023-3
Atlas	org.apache.atlas	sqoop-bridge-shim	3.0.0.7.1.9.1023-3
Avro	org.apache.avro	avro	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1023-3
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1023-3
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1023-3
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1023-3
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1023-3

Project	groupId	artifactId	version
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1023-3
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1023-3

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1023-3
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1023-3
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1023-3

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1023-3
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1023-3
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1023-3

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1023-3
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1023-3
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1023-3
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1023-3

Project	groupId	artifactId	version
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1023-3

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1023-3
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1023-3
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1023-3

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-knoxsso	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-knoxssout	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1023-3

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1023-3
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1023-3
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1023-3
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1023-3
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1023-3
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1023-3
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1023-3
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1023-3
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1023-3
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1023-3
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1023-3
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1023-3
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1023-3
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1023-3
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1023-3
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1023-3

Project	groupId	artifactId	version
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1023-3
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1023-3
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1023-3
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1023-3
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1023-3
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1023-3
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1023-3
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1023-3
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-analyzers-openslp	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1023-3
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1023-3

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1023-3
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1023-3
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1023-3
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1023-3
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1023-3
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1023-3
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-erasurencode	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-managed-rocksdbs	1.4.0.7.1.9.1023-3

Project	groupId	artifactId	version
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-filessystem	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-filessystem-common	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-filessystem-hadoop2	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-filessystem-hadoop3	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-filessystem-shaded	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-httpfsgateway	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1023-3
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1023-3

Project	groupId	artifactId	version
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1023-3
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compatible	6.0.0.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.1.6	5.1.1.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.2.5	5.1.1.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.3.0	5.1.1.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.0	5.1.1.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.1	5.1.1.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1023-3
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1023-3
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1023-3

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1023-3

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1023-3
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1023-3
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1023-3
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1023-3
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1023-3

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1023-3
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1023-3
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1023-3
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1023-3
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1023-3

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1023-3
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1023-3
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1023-3
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1023-3
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1023-3
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1023-3
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1023-3
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1023-3
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1023-3
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1023-3
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1023-3
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1023-3
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1023-3
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1023-3
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1023-3
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1023-3
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.1.9.1023-3
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1023-3
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1023-3
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1023-3
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1023-3
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1023-3
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1023-3

Runtime 7.1.9.1015-6

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	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-aws-s3-bridge	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-azure-adls-bridge	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-classification-updater	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-client-common	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-client-v1	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-client-v2	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-client-v2-shaded	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-common	3.0.0.7.1.9.1015-6

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-distro	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-docs	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-graphdb-api	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-graphdb-common	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-graphdb-janus	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-hdfs-bridge	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-index-repair-tool	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-intg	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-notification	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-plugin-classloader	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-repository	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-server-api	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	atlas-testtools	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	hbase-bridge	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	hbase-bridge-shim	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	hbase-testing-util	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	hdfs-model	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	hive-bridge	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	hive-bridge-shim	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	impala-bridge	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	impala-bridge-shim	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	impala-hook-api	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	kafka-bridge	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	kafka-bridge-shim	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	navigator-to-atlas	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	sample-app	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	sqoop-bridge	3.0.0.7.1.9.1015-6
Atlas	org.apache.atlas	sqoop-bridge-shim	3.0.0.7.1.9.1015-6
Avro	org.apache.avro	avro	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1015-6

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1015-6
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1015-6
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1015-6
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1015-6
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1015-6
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1015-6
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1015-6

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1015-6

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1015-6
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1015-6
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1015-6

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1015-6
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1015-6
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1015-6

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1015-6
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1015-6
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1015-6
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1015-6

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1015-6
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1015-6
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1015-6

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-knoxsso	1.3.0.7.1.9.1015-6

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-knoxssout	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1015-6
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1015-6
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1015-6
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1015-6
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1015-6

Project	groupId	artifactId	version
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1015-6
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1015-6
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1015-6
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1015-6
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1015-6
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1015-6
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1015-6
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1015-6
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1015-6
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1015-6
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1015-6
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1015-6
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1015-6
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1015-6
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1015-6
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1015-6
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1015-6
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1015-6
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1015-6
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-analyzers-opennlp	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1015-6

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1015-6
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1015-6
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1015-6
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1015-6
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1015-6
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1015-6
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1015-6
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1015-6

Project	groupId	artifactId	version
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-filessystem	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-filessystem-common	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-filessystem-hadoop2	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-filessystem-hadoop3	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-filessystem-shaded	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-httpsgateway	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1015-6
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1015-6

Project	groupId	artifactId	version
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1015-6
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compatible	6.0.0.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.1.6	5.1.1.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.2.5	5.1.1.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.3.0	5.1.1.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.0	5.1.1.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-hbase-compatible-2.4.1	5.1.1.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1015-6
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1015-6

Project	groupId	artifactId	version
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1015-6

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1015-6
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1015-6
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1015-6
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1015-6
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1015-6
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1015-6
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1015-6
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1015-6
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1015-6

Project	groupId	artifactId	version
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1015-6
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1015-6
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1015-6
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1015-6
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1015-6
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1015-6
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1015-6
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1015-6
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1015-6
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1015-6
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1015-6
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1015-6
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1015-6
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1015-6

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1015-6
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1015-6
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1015-6
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1015-6
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1015-6
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1015-6
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1015-6
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1015-6
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1015-6
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1015-6
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1015-6
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1015-6
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1015-6
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1015-6
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1015-6
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1015-6
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.1.9.1015-6
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1015-6
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1015-6
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1015-6
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1015-6
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1015-6
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1015-6

Runtime 7.1.9.1010-2

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	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-aws-s3-bridge	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-azure-adls-bridge	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-classification-updater	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-client-common	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-client-v1	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-client-v2	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-client-v2-shaded	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-common	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-distro	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-docs	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-graphdb-api	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-graphdb-common	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-graphdb-janus	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-hdfs-bridge	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-index-repair-tool	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-intg	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-notification	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-plugin-classloader	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-repository	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-server-api	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	atlas-testtools	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	hbase-bridge	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	hbase-bridge-shim	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	hbase-testing-util	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	hdfs-model	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	hive-bridge	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	hive-bridge-shim	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	impala-bridge	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	impala-bridge-shim	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	impala-hook-api	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	kafka-bridge	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	kafka-bridge-shim	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	navigator-to-atlas	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	sample-app	3.0.0.7.1.9.1010-2
Atlas	org.apache.atlas	sqoop-bridge	3.0.0.7.1.9.1010-2

Project	groupId	artifactId	version
Atlas	org.apache.atlas	sqoop-bridge-shim	3.0.0.7.1.9.1010-2
Avro	org.apache.avro	avro	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1010-2
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1010-2
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1010-2
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1010-2
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1010-2
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1010-2
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1010-2

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1010-2

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1010-2
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1010-2
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1010-2

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1010-2
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1010-2
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1010-2

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1010-2
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1010-2
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1010-2
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1010-2

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1010-2
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1010-2
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1010-2

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1010-2

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-knoxsso	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-knoxsout	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1010-2

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1010-2
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1010-2
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1010-2
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1010-2
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1010-2
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1010-2
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1010-2
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1010-2
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1010-2
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1010-2
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1010-2
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1010-2
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1010-2
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1010-2
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1010-2
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1010-2
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1010-2
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1010-2
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1010-2
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1010-2
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1010-2
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1010-2
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1010-2
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1010-2
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-analyzers-opennlp	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.1.9.1010-2

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1010-2
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1010-2

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1010-2
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1010-2
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1010-2
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1010-2
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1010-2
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1010-2
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-managed-rocksdb	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-filesystem	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-filesystem-common	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-filesystem-hadoop2	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-filesystem-hadoop3	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-filesystem-shaded	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-httpsgateway	1.4.0.7.1.9.1010-2

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1010-2
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1010-2
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compat	6.0.0.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.1.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.1.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.1.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.1.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.1.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1010-2

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1010-2
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1010-2
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1010-2

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1010-2

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1010-2
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1010-2
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1010-2
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1010-2
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1010-2

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1010-2
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1010-2
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1010-2
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1010-2
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1010-2
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1010-2
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1010-2
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1010-2
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1010-2
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1010-2
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1010-2
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1010-2
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1010-2
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1010-2
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1010-2
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1010-2
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1010-2

Project	groupId	artifactId	version
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1010-2
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1010-2
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1010-2
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.1.9.1010-2
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1010-2
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1010-2
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1010-2
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1010-2
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1010-2
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1010-2

Runtime 7.1.9.1005-2

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	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-aws-s3-bridge	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-azure-adls-bridge	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-classification-updater	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-client-common	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-client-v1	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-client-v2	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-client-v2-shaded	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-common	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-distro	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-docs	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-graphdb-api	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-graphdb-common	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-graphdb-janus	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-hdfs-bridge	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-index-repair-tool	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-intg	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-notification	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-plugin-classloader	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-repository	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-server-api	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	atlas-testtools	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	hbase-bridge	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	hbase-bridge-shim	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	hbase-testing-util	3.0.0.7.1.9.1005-2

Project	groupId	artifactId	version
Atlas	org.apache.atlas	hdfs-model	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	hive-bridge	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	hive-bridge-shim	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	impala-bridge	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	impala-bridge-shim	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	impala-hook-api	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	kafka-bridge	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	kafka-bridge-shim	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	navigator-to-atlas	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	sample-app	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	sqoop-bridge	3.0.0.7.1.9.1005-2
Atlas	org.apache.atlas	sqoop-bridge-shim	3.0.0.7.1.9.1005-2
Avro	org.apache.avro	avro	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-android	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-codegen-test	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-grpc	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-ipc-netty	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-maven-plugin	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-perf	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-protobuf	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-service-archetype	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-test-custom-conversions	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-thrift	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1005-2
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1005-2
Calcite	org.apache.calcite	calcite-babel	1.19.0.7.1.9.1005-2
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1005-2
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1005-2
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1005-2
Calcite	org.apache.calcite	calcite-server	1.19.0.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1005-2

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-assemblies	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-build-tools	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-client-api	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-client-miniclust	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-datajoin	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-extras	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-fs2img	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-gridmix	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-kafka	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-kms	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-miniclust	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1005-2

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-nfs	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-rumen	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-sls	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.1.9.1005-2
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1005-2
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-checkstyle	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-client-project	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-endpoint	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-examples	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-external-blockcache	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1005-2

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-hbtop	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-it	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-resource-bundle	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-rest	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-rsgroup	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-shaded-client-project	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-shaded-testing-util	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-shell	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-thrift	2.4.17.7.1.9.1005-2
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1005-2
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-contrib	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.1.9.1005-2

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-it-util	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-jmh	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-kryo-registrator	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-pre-upgrade	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-testutils	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1005-2
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1005-2
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1005-2
Kafka	org.apache.kafka	ci	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	connect	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	connect-basic-auth-extension	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	connect-cloudera-common	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	connect-file	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	connect-json	3.4.1.7.1.9.1005-2

Project	groupId	artifactId	version
Kafka	org.apache.kafka	connect-mirror	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	connect-mirror-client	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	connect-runtime	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	connect-transforms	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	generator	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-examples	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-shell	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-examples	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-test-utils	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.4.1.7.1.9.1005-2

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka-tools	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1005-2
Kafka	org.apache.kafka	trogdor	3.4.1.7.1.9.1005-2
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-docker	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-performance-test	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1005-2

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-idbroker	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-idbroker-plugins	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-knoxsso	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-knoxssout	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-vault	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-shell-release	1.3.0.7.1.9.1005-2

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-test	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-test-idbroker	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	webhdfs-kerb-test	1.3.0.7.1.9.1005-2
Knox	org.apache.knox	webhdfs-test	1.3.0.7.1.9.1005-2
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1005-2
Kudu	org.apache.kudu	kudu-backup2_2.11	1.17.0.7.1.9.1005-2
Kudu	org.apache.kudu	kudu-backup3_2.12	1.17.0.7.1.9.1005-2
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1005-2
Kudu	org.apache.kudu	kudu-hive	1.17.0.7.1.9.1005-2
Kudu	org.apache.kudu	kudu-spark2-tools_2.11	1.17.0.7.1.9.1005-2
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1005-2
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.17.0.7.1.9.1005-2
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1005-2
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1005-2
Livy	org.apache.livy	livy-api	0.7.2.7.1.9.1005-2
Livy	org.apache.livy	livy-client-common	0.7.2.7.1.9.1005-2
Livy	org.apache.livy	livy-client-http	0.7.2.7.1.9.1005-2
Livy	org.apache.livy	livy-core_2.11	0.7.2.7.1.9.1005-2
Livy	org.apache.livy	livy-examples	0.7.2.7.1.9.1005-2
Livy	org.apache.livy	livy-integration-test	0.7.2.7.1.9.1005-2
Livy	org.apache.livy	livy-repl_2.11	0.7.2.7.1.9.1005-2
Livy	org.apache.livy	livy-rsc	0.7.2.7.1.9.1005-2
Livy	org.apache.livy	livy-scala-api_2.11	0.7.2.7.1.9.1005-2
Livy	org.apache.livy	livy-server	0.7.2.7.1.9.1005-2

Project	groupId	artifactId	version
Livy	org.apache.livy	livy-test-lib	0.7.2.7.1.9.1005-2
Livy	org.apache.livy	livy-thriftserver	0.7.2.7.1.9.1005-2
Livy	org.apache.livy	livy-thriftserver-session	0.7.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-analyzers-openslp	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1005-2
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1005-2

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-server	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.1.9.1005-2
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.1.9.1005-2
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1005-2
ORC	org.apache.orc	orc-examples	1.5.1.7.1.9.1005-2
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1005-2
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1005-2
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-container-service	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-docs	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-hadoop-dependency-server	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-hadoop-dependency-test	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-interface-server	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-managed-rocksdbs	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-rocks-native	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-server-framework	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-server-scm	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-test-utils	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	hdds-tools	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	mini-chaos-tests	1.4.0.7.1.9.1005-2

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-csi	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-datanode	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-dist	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-filessystem	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-filessystem-common	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-filessystem-hadoop2	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-filessystem-hadoop3	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-filessystem-shaded	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-httpfsgateway	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-insight	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-integration-test	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-interface-storage	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-manager	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-network-tests	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-recon	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-reconcodegen	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-s3-secret-store	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-s3gateway	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	ozone-tools	1.4.0.7.1.9.1005-2
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	1.4.0.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-cascading	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-cascading3	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-generator	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-pig	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-pig-bundle	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-protobuf	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-scala_2.10	1.10.99.7.1.9.1005-2
Parquet	org.apache.parquet	parquet-thrift	1.10.99.7.1.9.1005-2

Project	groupId	artifactId	version
Parquet	org.apache.parquet	parquet-tools	1.10.99.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-connectors-phoenix5-compat	6.0.0.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-core	5.1.1.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.1.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.1.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.1.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.1.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.1.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-pherf	5.1.1.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix5-spark	6.0.0.7.1.9.1005-2
Phoenix	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.1.9.1005-2
Ranger	org.apache.ranger	conditions-enrichers	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	embeddedwebserver	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	jisql	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ldapconfigcheck	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-adls-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-atlas-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-distro	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-examples-distro	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-hbase-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-hive-plugin	2.4.0.7.1.9.1005-2

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-intg	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-kafka-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-kms-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-knox-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-kudu-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-kylin-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-nifi-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-ozone-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-plugins-installer	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-policymigration	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-raz-adls	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-rms-tools	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-solr-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.4.0.7.1.9.1005-2

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-storm-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-tools	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-yarn-plugin	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	sample-client	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	sampleapp	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	shaded-raz-hook-abfs	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1005-2
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1005-2
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-analytics	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-cell	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-langid	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-ltr	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-security-util	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-test-framework	8.11.2.7.1.9.1005-2
Solr	org.apache.solr	solr-velocity	8.11.2.7.1.9.1005-2
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1005-2

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1005-2
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1005-2
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1005-2
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.1.9.1005-2
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-aux-services	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-tests	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.1.9.1005-2
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1005-2

Project	groupId	artifactId	version
Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.1.9.1005-2
Zeppelin	org.apache.zeppelin	zeppelin-display	0.8.2.7.1.9.1005-2
Zeppelin	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.1.9.1005-2
Zeppelin	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.1.9.1005-2
Zeppelin	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.1.9.1005-2
Zeppelin	org.apache.zeppelin	zeppelin-livy	0.8.2.7.1.9.1005-2
Zeppelin	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.1.9.1005-2
Zeppelin	org.apache.zeppelin	zeppelin-server	0.8.2.7.1.9.1005-2
Zeppelin	org.apache.zeppelin	zeppelin-shell	0.8.2.7.1.9.1005-2
Zeppelin	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.1.9.1005-2
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1005-2
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.1.7.1.9.1005-2
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.1.7.1.9.1005-2
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.1.7.1.9.1005-2
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.1.7.1.9.1005-2
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.1.7.1.9.1005-2
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1005-2
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.1.7.1.9.1005-2
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.1.7.1.9.1005-2
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.1.7.1.9.1005-2
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.1.7.1.9.1005-2

Runtime 7.1.9.1000-103

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	3.0.0.7.1.9.1000-103
Atlas	org.apache.atlas	atlas-client-common	3.0.0.7.1.9.1000-103
Atlas	org.apache.atlas	atlas-client-v1	3.0.0.7.1.9.1000-103
Atlas	org.apache.atlas	atlas-client-v2	3.0.0.7.1.9.1000-103
Atlas	org.apache.atlas	atlas-client-v2-shaded	3.0.0.7.1.9.1000-103
Atlas	org.apache.atlas	atlas-common	3.0.0.7.1.9.1000-103
Atlas	org.apache.atlas	atlas-intg	3.0.0.7.1.9.1000-103
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	3.0.0.7.1.9.1000-103
Atlas	org.apache.atlas	atlas-notification	3.0.0.7.1.9.1000-103
Atlas	org.apache.atlas	atlas-plugin-classloader	3.0.0.7.1.9.1000-103
Atlas	org.apache.atlas	hive-bridge	3.0.0.7.1.9.1000-103
Avro	org.apache.avro	avro	1.11.1.7.1.9.1000-103
Avro	org.apache.avro	avro-compiler	1.11.1.7.1.9.1000-103
Avro	org.apache.avro	avro-ipc	1.11.1.7.1.9.1000-103
Avro	org.apache.avro	avro-ipc-jetty	1.11.1.7.1.9.1000-103

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-mapred	1.11.1.7.1.9.1000-103
Avro	org.apache.avro	avro-tools	1.11.1.7.1.9.1000-103
Avro	org.apache.avro	trevni-avro	1.11.1.7.1.9.1000-103
Avro	org.apache.avro	trevni-core	1.11.1.7.1.9.1000-103
Calcite	org.apache.calcite	calcite-core	1.19.0.7.1.9.1000-103
Calcite	org.apache.calcite	calcite-druid	1.19.0.7.1.9.1000-103
Calcite	org.apache.calcite	calcite-linq4j	1.19.0.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-annotations	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-archives	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-auth	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-aws	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-azure	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-client	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-common	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-distcp	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-hdfs	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-minicluster	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-minikdc	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-openstack	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-streaming	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.1.9.1000-103

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.1.9.1000-103
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.1.9.1000-103
HBase	org.apache.hbase	hbase-annotations	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-asyncfs	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-client	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-common	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-hadoop-compat	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-hadoop2-compat	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-http	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-logging	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-mapreduce	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-metrics	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-metrics-api	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-procedure	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-protocol	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-protocol-shaded	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-replication	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-server	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-shaded-client	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-testing-util	2.4.17.7.1.9.1000-103
HBase	org.apache.hbase	hbase-zookeeper	2.4.17.7.1.9.1000-103
Hive	org.apache.hive	hive-beeline	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-classification	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-cli	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-common	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-exec	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-metastore	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-serde	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-service	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-shims	3.1.3000.7.1.9.1000-103

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-spark-client	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-streaming	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.1.9.1000-103
Hive	org.apache.hive	kafka-handler	3.1.3000.7.1.9.1000-103
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.1.9.1000-103
Kafka	org.apache.kafka	connect-api	3.4.1.7.1.9.1000-103
Kafka	org.apache.kafka	kafka-clients	3.4.1.7.1.9.1000-103
Kafka	org.apache.kafka	kafka-group-coordinator	3.4.1.7.1.9.1000-103
Kafka	org.apache.kafka	kafka-log4j-appender	3.4.1.7.1.9.1000-103
Kafka	org.apache.kafka	kafka-metadata	3.4.1.7.1.9.1000-103
Kafka	org.apache.kafka	kafka-raft	3.4.1.7.1.9.1000-103
Kafka	org.apache.kafka	kafka-server-common	3.4.1.7.1.9.1000-103
Kafka	org.apache.kafka	kafka-storage	3.4.1.7.1.9.1000-103
Kafka	org.apache.kafka	kafka-storage-api	3.4.1.7.1.9.1000-103
Kafka	org.apache.kafka	kafka-streams	3.4.1.7.1.9.1000-103
Kafka	org.apache.kafka	kafka_2.12	3.4.1.7.1.9.1000-103
Kafka	org.apache.kafka	kafka_2.13	3.4.1.7.1.9.1000-103
Knox	org.apache.knox	gateway-adapter	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-admin-ui	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-applications	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-cloud-bindings	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-demo-ldap	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-discovery-ambari	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-discovery-cm	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-i18n	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-ha	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-jersey	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-rewrite	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.1.9.1000-103

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-release	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-server	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-server-launcher	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-admin	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-as	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-definitions	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-hbase	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-health	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-hive	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-impala	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-jkg	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-knoxsso	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-knoxsout	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-livy	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-metadata	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-nifi	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-rm	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-session	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-storm	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-test	1.3.0.7.1.9.1000-103

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-tgs	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-service-webhdfs	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-shell	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-shell-launcher	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-shell-samples	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-spi	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-test-release-utils	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-test-utils	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-topology-simple	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-util-common	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-util-configinjector	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-util-launcher	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	gateway-util-urltemplate	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	hadoop-examples	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	knox-cli-launcher	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	knox-homepage-ui	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	knox-token-generation-ui	1.3.0.7.1.9.1000-103
Knox	org.apache.knox	knox-token-management-ui	1.3.0.7.1.9.1000-103
Kudu	org.apache.kudu	kudu-backup-tools	1.17.0.7.1.9.1000-103
Kudu	org.apache.kudu	kudu-client	1.17.0.7.1.9.1000-103
Kudu	org.apache.kudu	kudu-spark2_2.11	1.17.0.7.1.9.1000-103
Kudu	org.apache.kudu	kudu-spark3_2.12	1.17.0.7.1.9.1000-103
Kudu	org.apache.kudu	kudu-test-utils	1.17.0.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-core	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-join	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.1.9.1000-103

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.1.9.1000-103
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.1.9.1000-103
Oozie	org.apache.oozie	oozie-client	5.1.0.7.1.9.1000-103
Oozie	org.apache.oozie	oozie-core	5.1.0.7.1.9.1000-103
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.1.9.1000-103
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.1.9.1000-103
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.1.9.1000-103
Oozie	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.1.9.1000-103
ORC	org.apache.orc	orc-core	1.5.1.7.1.9.1000-103
ORC	org.apache.orc	orc-mapreduce	1.5.1.7.1.9.1000-103
ORC	org.apache.orc	orc-shims	1.5.1.7.1.9.1000-103
ORC	org.apache.orc	orc-tools	1.5.1.7.1.9.1000-103
Ozone	org.apache.ozone	hdds-annotation-processing	1.4.0.7.1.9.1000-103
Ozone	org.apache.ozone	hdds-client	1.4.0.7.1.9.1000-103
Ozone	org.apache.ozone	hdds-common	1.4.0.7.1.9.1000-103
Ozone	org.apache.ozone	hdds-config	1.4.0.7.1.9.1000-103
Ozone	org.apache.ozone	hdds-erasurecode	1.4.0.7.1.9.1000-103
Ozone	org.apache.ozone	hdds-hadoop-dependency-client	1.4.0.7.1.9.1000-103
Ozone	org.apache.ozone	hdds-interface-admin	1.4.0.7.1.9.1000-103
Ozone	org.apache.ozone	hdds-interface-client	1.4.0.7.1.9.1000-103
Ozone	org.apache.ozone	ozone-client	1.4.0.7.1.9.1000-103
Ozone	org.apache.ozone	ozone-common	1.4.0.7.1.9.1000-103
Ozone	org.apache.ozone	ozone-filessystem	1.4.0.7.1.9.1000-103
Ozone	org.apache.ozone	ozone-filessystem-common	1.4.0.7.1.9.1000-103
Ozone	org.apache.ozone	ozone-filessystem-hadoop3	1.4.0.7.1.9.1000-103
Ozone	org.apache.ozone	ozone-filessystem-shaded	1.4.0.7.1.9.1000-103
Ozone	org.apache.ozone	ozone-interface-client	1.4.0.7.1.9.1000-103
Parquet	org.apache.parquet	parquet-avro	1.10.99.7.1.9.1000-103
Parquet	org.apache.parquet	parquet-column	1.10.99.7.1.9.1000-103
Parquet	org.apache.parquet	parquet-common	1.10.99.7.1.9.1000-103
Parquet	org.apache.parquet	parquet-encoding	1.10.99.7.1.9.1000-103
Parquet	org.apache.parquet	parquet-format-structures	1.10.99.7.1.9.1000-103
Parquet	org.apache.parquet	parquet-hadoop	1.10.99.7.1.9.1000-103
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.1.9.1000-103
Parquet	org.apache.parquet	parquet-jackson	1.10.99.7.1.9.1000-103
Phoenix	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.1.9.1000-103

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.1.9.1000-103
Ranger	org.apache.ranger	credentialbuilder	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-authn	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-common-ha	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-kms	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-metrics	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-plugin-classloader	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-plugins-audit	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-plugins-common	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-plugins-cred	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-raz-intg	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-raz-processor	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-rms-common	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-rms-hive	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-rms-webapp	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-tagsync	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ranger-util	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	ugsync-util	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	unixauthclient	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	unixauthservice	2.4.0.7.1.9.1000-103
Ranger	org.apache.ranger	unixusersync	2.4.0.7.1.9.1000-103
Solr	org.apache.solr	solr-core	8.11.2.7.1.9.1000-103
Solr	org.apache.solr	solr-solrj	8.11.2.7.1.9.1000-103
Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-catalyst_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-core_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-graphx_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-hive_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-kvstore_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-launcher_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-mllib_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-network-common_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.1.9.1000-103

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-repl_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-sketch_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-sql_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-streaming_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-tags_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-unsafe_2.11	2.4.8.7.1.9.1000-103
Spark	org.apache.spark	spark-yarn_2.11	2.4.8.7.1.9.1000-103
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.1.9.1000-103
Tez	org.apache.tez	hadoop-shim	0.9.1.7.1.9.1000-103
Tez	org.apache.tez	tez-api	0.9.1.7.1.9.1000-103
Tez	org.apache.tez	tez-common	0.9.1.7.1.9.1000-103
Tez	org.apache.tez	tez-dag	0.9.1.7.1.9.1000-103
Tez	org.apache.tez	tez-examples	0.9.1.7.1.9.1000-103
Tez	org.apache.tez	tez-history-parser	0.9.1.7.1.9.1000-103
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.1.9.1000-103
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.1.9.1000-103
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.1.9.1000-103
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.1.9.1000-103
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.1.9.1000-103
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.1.9.1000-103
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.1.7.1.9.1000-103
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.1.7.1.9.1000-103

Fixed issues in Cloudera Runtime 7.1.9 SP1

You can review the list of reported issues and their fixes in Cloudera Runtime 7.1.9 SP1. Fixed issues represent selected issues that were previously logged through Cloudera Support, but are now addressed in the current Runtime release. These issues may have been reported in previous versions of Runtime as a known issue; meaning they were reported by customers or identified by Cloudera Quality Engineering teams.

Fixed Issues in Atlas

Review the list of Atlas issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-69894: Atlas uses the upstream iceberg-hive-runtime package instead of the CDP build

The version of the iceberg-hive-runtime artifact is changed to a snapshot version, so downstream patches, bugfixes and CVE fixes will be pulled.

CDPD-69162: Upgrade jersey version to jersey 2

The Jersey version is upgraded to use Jersey 2 after Jackson Databind is updated to 2.15.

CDPD-66567: fetchType as "incremental" does full export instead of "CONNECTED"

An additional check is added for the Hive and Iceberg table export request to check if the fetch type is incremental or not.

CDPD-65509: Atlas Import/Export with iceberg_table type is copying Hive External Entity also

An additional check is added for the Hive and Iceberg table export request to check if the fetch type is incremental or not.

CDPD-59028, CDPD-62893: Cloudera Runtime 7.1.9 SPx Audit aging should be possible with only ttl or audit count instead of both

The issue was fixed by enabling the separate configuration of time-to-live and audit count limit parameters. For more information, see [Using default audit aging](#)

CDPD-62837: hive_storagedesc is incomplete in Atlas

The HDFS data pertaining to Hive tables that have undergone migration using the Navigator to Atlas transition process might be missing some information. The Nav2Atlas module was changed to set the relationType as "hive_table_storagedesc" of relationship attribute "table" for every entity of "hive_storagedesc".

CDPD-62834: Status of deleted table is ACTIVE in Atlas after navigator2atlas migration

Deleted Hive tables migrated using the Navigator to Atlas transition may be shown as active in Atlas. The Nav2Atlas module was changed to set the relationType as "hive_table_storagedesc" of relationship attribute "table" for every entity of "hive_storagedesc".

CDPD-59033: Compression algorithm changed from Gzip to Snappy in Atlas HBase tables

The default compression algorithm was changed from Gzip to 'Snappy' in Atlas HBase tables.

CDPD-57021: Relationship Search feature should be configurable

Relationship Search feature is made configurable to reduce time taken for iterating all the edges.

CDPD-55697: Clean Atlas application log in case of HA mode

The unwanted error logs are removed from the Atlas application log file.

CDPD-67977: Nullpointer exception during incremental Atlas Export

An incremental export with glossary and terms associated with exported entities was causing a null pointer exception, which is now resolved.

CDPD-67553: Parallel Atlas Import is failing with unexpected errors

Import operation of an entity is retried thrice when the transaction fails because of a PermanentLockingException. This issue is now resolved.

CDPD-59586: OperationType does not support Contains operation

Support for sub-string and wild character '*' is not present for attributeName=operationType in the Custom Audit Filters feature. Hence, attributeValue can not be checked with the contains operator. This is fixed by adding support for substrings and the wildcard character "*" for the attribute operationType.

CDPD-65804: Atlas cannot search comments containing Chinese or Japanese characters

Entities whose comments contain Chinese or Japanese characters can be now searched by using the comment (string), containing filters.

CDPD-69654: Atlas parallel export fails with the formation of incomplete .zip files

An extra flag is added in Atlas. By enabling it (skipUpdateReplicationAttr:true), it will skip updating each entity for adding replicatedTo/replicatedFrom attribute, which will ultimately increase the performance of exporting and importing.

CDPD-69176: Entity search with "begins with" operator does not return correct results for some Japanese characters

Entities containing Japanese or Chinese can also be searched by using the "begins with" search filter.

CDPD-68655: Tag searches with multiple Japanese or multiple Chinese characters do not return the correct results

Tag searches work with multiple Japanese or Chinese characters included for `hdfs_path` entities.

CDPD-68643: Basic search query with '<Japanese search string>*' does not always return results

Searches including Japanese characters and the asterisk character can be reliably made.

CDPD-68328: Faceted `hdfs_path` entity type searches containing Japanese or Chinese characters do not return results only with query format searches.

`Hdfs_path` type entity queries can employ multiple filters (attribute, classification) concurrently.

CDPD-68189: Faceted searches do not give correct results for the "contains" operator applied to Chinese words with multiple characters.

Chinese words with multiple characters can be searched reliably.

CDPD-67022: Imports fail with 403 error by users with export, import permissions but without create, read, write, and update entity permissions

After the update, following a successful import, Atlas creates an entity of type `_{}AtlasAuditEntry{}` and no longer throws an `_UNAUTHORIZED_ACCESS` error with status code 403 because lacking create, read, write or update user permissions.

CDPD-67020: Imports fail by users with export or import permissions but no other permission for entities

After the update, during an export process, the following applies:

- The `AtlasServerService` creates a new `AtlasServer` entity if it does not exist already.
- The `xportImportAuditService` creates an entity with type `__ExportImportAuditEntry` if adequate read and write permissions are not present for the mentioned entities, exceptions are no longer thrown and the creation of the `atlas-export-result.json` is no longer skipped in the flow. The import with this export no longer throws a `NullPointerException` due to the missing `atlas-export-result.json`.

CDPD-67022: Imports fail with 403 error by users with export, import permissions but without create, read, write, and update entity permissions

After the update, following a successful import, Atlas creates an entity of type `_{}AtlasAuditEntry{}` and no longer throws an `_UNAUTHORIZED_ACCESS` error with status code 403 because lacking create, read, write or update user permissions.

Apache Jira: ATLAS-4836

Apache patch information

- ATLAS-4855
- ATLAS-4819
- ATLAS-4836
- ATLAS-4837

Fixed Issues in Arrow

Review the list of Arrow issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-69641: Use Centralised Guava version in Arrow

Upgraded guava library from version 28.1-jre to 32.0.1-jre due to CVE-2023-2976

Apache patch information**Fixed Issues in Avro**

Review the list of Avro issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-66671: Avro - Upgrade commons-compress to 1.26.0 due to CVE-2023-42503, CVE-2024-25710 and CVE-2024-26308

Upgraded commons-compress library from 1.22 to 1.26.0 due to CVE-2023-42503, CVE-2024-25710 & CVE-2024-26308

CDPD-63092: Out of memory issues due to CVE-2023-39410

When deserializing untrusted or corrupted data, it is possible for a reader to consume memory beyond the allowed constraints and therefore leading to out of memory on the system. This issue affects Java applications using Apache Avro Java SDK up to and including 1.11.2. Users should update to apache-avro version 1.11.3 which addresses this issue.

Fixed issues in Calcite

Review the list of Calcite issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-69166: Backport CALCITE-6280 to 7.1.9 SP1

Backport of CALCITE-6280

Apache Patch Information

- CALCITE-6280

Fixed issues in Cloud Connectors

There are no fixed issues for Cloud Connectors in Cloudera Runtime 7.1.9 SP1.

Fixed issues in Cruise Control

Review the list of Cruise Control issues that are resolved in Cloudera Runtime 7.1.9 SP1.

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

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

Fixed Issues in Hadoop

There are no fixed issues for Hadoop in Cloudera Runtime 7.1.9 SP1.

Apache Patch Information

None

Fixed Issues in HDFS

There are no fixed issues for HDFS in Cloudera Runtime 7.1.9 SP1.

CDPD-49723:CacheReplicationMonitor holds lock for extended period causing NameNode outage

CacheReplicationMonitor periodically scans cache directives and cached BlockMap. As more cache directives are added, the scan takes longer, holding a global write lock and preventing the NameNode from processing other requests. The issue is fixed in HDFS-15273, which addresses lock optimization for the effective utilization of Centralized Cache Management.

Apache Patch Information

None

Fixed Issues in HBase

Review the list of HBase issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-70354: Backport of HBASE-28595

This is a backport of HBASE-28595 - Losing exception from scan RPC can lead to partial results.

CDPD-69218: Backport of HBASE-27230

This is a backport of HBASE-27230 - RegionServer must be aborted when WAL . sync throws TimeoutIOException.

Apache Patch Information

- HBASE-28595
- HBASE-27230

Fixed Issues in Hive

Review the list of Hive issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-70428: Issue with legacy timestamp conversion in Parquet files

When converting legacy timestamps in Parquet files, the date 'February 29, 200' can cause issues. This happens because the Julian day for February 29, 200 CE, differs between the Julian and Gregorian calendars.

Hive, which stores timestamps in UTC, encounters this issue when converting the date '200 CE/03/01' between timezones. Even if the original date was '200 CE/03/01' in the Asia/Singapore timezone, the conversion process leads to the "not a leap year exception" for February 29, 200.

```
java.time.DateTimeException: Invalid date 'February 29' as '200'
is not a leap year
```

This issue has been fixed.

CDPD-67820: Select query failure due to decimal column data type change in Parquet

When you change the data type of a decimal column to STRING, CHAR, or VARCHAR in a Parquet file, the select queries on that column fail. This issue arises because the Parquet format handles these data types differently, causing mismatches and query failures.

This issue has been fixed.

CDPD-67819: Incorrect results for IN UDF on Parquet column of CHAR/VARCHAR type

When you use the IN UDF on a Parquet column of CHAR or VARCHAR type, you encounter incorrect results. To fix this issue CAST operations are added during the conversion process to properly handle these data types.

CDPD-33110: Hive queries stuck because of excessive aborted transactions

Hive queries can become unresponsive due to a high number of aborted transactions. The excessive number of aborted transactions causes the cleaner thread to consume all the memory of the Hive Metastore (HMS) while attempting to clean these transactions.

The issue was addressed by deleting aborted transactions directly instead of preloading them. This approach prevents the cleaner thread from overloading the HMS memory.

CDPD-69846: HPLSQL not using Hive variables

This issue is observed in CDP Private Cloud Base 7.1.9 and higher versions. When passing Hive variables using the --hivevar variable through Beeline, the variables are not utilized within HPLSQL procedures or cursors, leading to an unhandled exception in HPL/SQL.

This issue has been fixed.

CDPD-67350: DirectSQL and JDO results are different when fetching partitions by timestamp in DST shift

Fetching partition metadata from Hive Metastore using Java Data Objects (JDO) by timestamp does not provide correct result in DST shift in partition pruning.

The issue has been fixed.

CDPD-66902: Query runtime optimization

Optimized SQL query to remove duplicate records from COMPLETED_TXN_COMPONENTS to improve runtime.

CDPD-61611: Impala stats blocks Hive partitioned table rename

The issue occurs when the Impala tries to rename a Hive acid table, which has the Impala stats beforehand.

```
ERROR : Failed
org.apache.hadoop.hive.ql.metadata.HiveException: Unable to alter
table. Cannot change stats state for a transactional table defa
ult.parqtest without providing the transactional write state for
verification (new write ID 6, valid write IDs null; current sta
te null; new state {})
```

To address the issue, as the Impala rename does not change the table's stats, even though the write ID list for this table is absent, the absence is ignored when verifying the stats in such a case.

CDPD-55916: Confusing exception message in DagUtils.localizeResource

Hive jobs occasionally fail, during resource localization when a Tez session is created. DagUtils#localizeResource is responsible for copying the client's hive-exec.jar into HDFS. This process can be triggered from multiple threads concurrently, in which case one thread performs the copy while the others wait, polling for arrival of the destination file.

If there is an IOException during this process, it is assumed that the thread attempting the write failed, and all others abort. No information about the underlying IOException is logged. Instead, the log states:

```
java.io.IOException: Previous writer likely failed to write hdfs
:///....Failing because I am unlikely to write too
```

To address this issue, the logging is improved by stating that a failure on the writing thread is just one possible reason for the error. It also logs the exception stack trace to make it easier to find the real root cause.

CDPD-68278: Netty HttpPostRequestDecoder Vulnerability

Netty's HttpPostRequestDecoder can be tricked into accumulating unlimited data, leading to potential vulnerabilities.

To address the issue, upgraded Netty from version 4.1.100.Final to 4.1.108.Final.

CDPD-68251: Upgrade Commons-configuration2 to Version 2.10.1

To address the issue, upgrade Commons-configuration2 version to 2.10.1 to fix a crash caused by complex data structures.

Apache patch information

- HIVE-28249
- HIVE-26955
- HIVE-26320
- HIVE-27775
- HIVE-13288
- HIVE-28214
- HIVE-27778

Fixed Issues in Hue

Review the list of Hue issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-61651: Schema and tables are not being displayed in the Phoenix Editor as Phoenix is case-sensitive

Schema and table names in Phoenix are case-sensitive. Hue used the UPPER() function earlier while processing queries, resulting in issues when running Phoenix queries. This issue has been resolved by removing the UPPER() function for databases and tables.

CDPD-61550: ABFS File Browser lists a maximum of 5000 files

Earlier, when you browsed ADLS Gen 2 storage using Hue's ABFS File Browser, it displayed only 5000 objects. This issue has been fixed by including a continuation token in the response header (x-ms-continuation).

CDPD-45438: Hue Importer does not validate the column name; displays an error if the column name contains a space

If column names contained spaces in the file you were trying to import into Hue, then you encountered the following error: AnalysisException: Invalid column/field name. When you fixed the column name and proceeded with importing the file, the rename operation also failed. This issue has been fixed. Hue now validates the column names to verify whether they contain spaces and does not allow you to proceed without correcting the column names.



Note: Hue only allows alphanumeric characters and underscores (_) in column names for all dialects and engines.

CDPD-59802: Hue is unable to process parallel requests

An issue with the Impala file reader not getting closed after use caused a leak in the connection pool. Because of the leak, the Hue Query Processor could not read Impala event files after a few queries. This issue has been fixed by ensuring that the file reader is closed and that the ingestion of the Impala query works as expected.

Fixed Issues in Impala

Review the list of Impala issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-67912: Multiple fixes address performance issues with improved memory tracking for row batch serialization during KRPC shuffling

Queries shuffling large row batches were consuming untracked memory during row batch serialization, leading to higher memory usage without proper accounting.

Memory for tuple_data_ and compression scratch is now tracked using MemTrackerAllocator, and memory estimation. Added to account for outbound row batch allocation. Memory for row batch serialization is now fully tracked, leading to more accurate memory usage in queries that shuffle rows.

CDPD-67901: Support python 3.10

The impala-shell commands now works with Python 3.10 on SLES 15 operating system.

CDPD-60005: Fix for two Catalogd with force active failure

The issue has been resolved.

CDPD-67493: Catalogd performance fix

The Catalogd performance is enhanced by the following improvements done through:

- Impala now reloads events to update metadata across Catalogd replicas for refresh or invalidate commands.
- Impala now handles events thrown by Hive Metastore when a compaction is committed, allowing for automatic refreshing after compactions.
- Impala now skips redundant refresh events by processing only the first event, reducing overhead.
- Catalogd now skips preceding events after a manual table refresh, allowing it to quickly catch up with Hive Metastore events.
- Impala now efficiently processes notification events in Catalogd.
- Refined batching logic in Catalogd optimizes event processing by grouping events across different sources per table, enhancing efficiency.

- Impala now optimizes event processing by skipping unnecessary events and better handling manual refreshes to prevent disruptions.
- Catalogd's event processor now reloads file metadata specifically for truncate table operations, in Hive Metastore alter table events.
- Impala's event batching now efficiently manages events for a table, even if there are unrelated events in between.
- Impala now includes table-level metrics added to the Catalogd webUI page, such as avg-events-process-duration and events-consuming-delay.
- Impala now accurately processes subsequent `ALTER_PARTITION` events after creating a partition through `INSERT` operation, ensuring enhanced event processing stability and accuracy.

Apache Patch Information

- IMPALA-12123
- IMPALA-10120
- IMPALA-10987
- IMPALA-11284
- IMPALA-11710
- IMPALA-12827
- IMPALA-11805
- IMPALA-11901
- IMPALA-11909
- IMPALA-11918
- IMPALA-11921
- IMPALA-11924
- IMPALA-12001
- IMPALA-12002
- IMPALA-12049
- IMPALA-12082
- IMPALA-12131
- IMPALA-12213
- IMPALA-12233
- IMPALA-12256
- IMPALA-12257
- IMPALA-12260
- IMPALA-12267
- IMPALA-12269
- IMPALA-12273
- IMPALA-12274
- IMPALA-12292
- IMPALA-12299
- IMPALA-12306
- IMPALA-12311
- IMPALA-12322
- IMPALA-12327
- IMPALA-12228
- IMPALA-12333
- IMPALA-12340
- IMPALA-12341
- IMPALA-12342
- IMPALA-12359

- IMPALA-12366
- IMPALA-12371
- IMPALA-12383
- IMPALA-12393
- IMPALA-12395
- IMPALA-12397
- IMPALA-12403
- IMPALA-12416
- IMPALA-12417
- IMPALA-12429
- IMPALA-12444
- IMPALA-12464
- IMPALA-12478
- IMPALA-12508
- IMPALA-12510
- IMPALA-12511
- IMPALA-12513
- IMPALA-12517
- IMPALA-12524
- IMPALA-12525
- IMPALA-12526
- IMPALA-12532
- IMPALA-12535
- IMPALA-12545
- IMPALA-12550
- IMPALA-12552
- IMPALA-12555
- IMPALA-12556
- IMPALA-12561
- IMPALA-12565
- IMPALA-12579
- IMPALA-12581
- IMPALA-12614
- IMPALA-12617
- IMPALA-12630
- IMPALA-12632
- IMPALA-12633
- IMPALA-12643
- IMPALA-12666
- IMPALA-12681
- IMPALA-12698
- IMPALA-12721
- IMPALA-6590
- IMPALA-6876
- IMPALA-12894
- IMPALA-11430
- IMPALA-12833
- IMPALA-12849
- IMPALA-12529
- IMPALA-12711

- IMPALA-12832
- IMPALA-12831
- IMPALA-12929
- IMPALA-12699
- IMPALA-12356
- IMPALA-12719
- IMPALA-12463
- IMPALA-11768
- IMPALA-11553
- IMPALA-12356
- IMPALA-10949
- IMPALA-12399
- IMPALA-12356
- IMPALA-12719
- IMPALA-12463
- IMPALA-6684

Fixed Issues in Iceberg

Review the list of Iceberg issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-45139: Improve Iceberg V2 reads with a custom Iceberg Position Delete operator

This fix helps by improving the Impala query performance when reading Iceberg tables with delete files.

CDPD-47349: Use hive.metastore.table.owner during table creation

Creation of an Iceberg table used to happen in two steps. The first step creates a table, however, the table is created with the wrong owner. The second step runs the ALTER TABLE statement to set the correct table owner.

This fix resolves the issue by creating an Iceberg table with the correct owner in a single step. When creating an Iceberg table, Impala specifies the owner in the hive.metastore.table.owner table property.

CDPD-55029: Include snapshot ID of Iceberg tables in query plan or profile

This fix includes the snapshot ID of Iceberg tables for the Iceberg SCAN operators, which is useful to enable queries to be re-executed. Re-executing queries are useful because they help you to better investigate performance problems and bugs.

CDPD-59657: Iceberg V2 operator provides incorrect results in PARTITIONED mode

When PARTITIONED mode is used, the fix introduced through [IMPALA-12327](#) performs a binary search when the position-based difference between the current row and previous row is not one.

CDPD-60282: Need better cardinality estimation for Iceberg V2 tables with deletes

Currently, the cardinality of the IcebergDeleteNode is the same as the cardinality of the left-hand side (LHS) and does not take into account the cardinality of the right-hand side (RHS). The RHS contains position delete records, therefore, all the records in RHS remove a record from the LHS.

If there are joins on the Iceberg table, they have the same selectivity on the data records and on the delete records.

This fix updates the cardinality of the IcebergDeleteNode to use the following formula:

$$\text{Cardinality of DELETE operator} = \text{Cardinality(LHS)} - (\text{Cardinality(RHS)} * \text{selectivity of LHS})$$

CDPD-60946/CDPD-60717: Iceberg tables created through Trino are incompatible with Impala

The Trino SQL engine creates Iceberg tables without setting `hive.engine.enabled=true` and does not provide users with an option to manually set this property. Therefore, Trino always creates Iceberg tables with non-HiveIceberg storage descriptors.

Impala uses the Input/Output/SerDe properties to determine the table type, however, a table is also considered to be an Iceberg table if the table property, `table_type=ICEBERG` is set.

The fix introduced through [IMPALA-12413](#) ensures that modifications to the table from Impala goes through its Iceberg library (with `engine.hive.enabled=true`). This results in setting the HiveIceberg storage descriptors and allows Trino to be compatible with Iceberg tables.

CDPD-66786: Impala returns incorrect results when the optimized Iceberg V2 operator is used

If you are using Impala to read Iceberg V2 tables, then you might have noticed Impala returning incorrect results when the optimized V2 operator is used. This issue has been resolved by resetting the delete state when it detects records from files that do not have delete records.

CDPD-67632: Optimized count(*) for Iceberg table gives wrong results after a Spark rewrite_data_files

Spark's `rewrite_data_files` action can leave dangling Delete files in the Iceberg table. Delete files are not applicable to any data files. This can cause incorrect results in Impala when it runs simple `count(*)` queries with the help of table statistics stored in the Iceberg metadata layer.

This fix resolves the issue and Impala now returns the correct results even with the presence of dangling Delete files.

Apache Patch Information

- IMPALA-11619
- IMPALA-11776
- IMPALA-12072
- IMPALA-12327
- IMPALA-12371
- IMPALA-12413
- IMPALA-12894

Technical Service Bulletins

TSB 2024-752: Dangling delete issue in Spark rewrite_data_files procedure causes incorrect results for Iceberg V2 tables

For the latest update on this issue see the corresponding Knowledge article: [TSB 2024-752: Dangling delete issue in Spark rewrite_data_files procedure causes incorrect results for Iceberg V2 tables](#).

Fixed Issues in Kafka

Review the list of Kafka issues that are resolved in Cloudera Runtime 7.1.9 SP1.

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

Cloudera Manager now registers the metrics `kafka_connect_connector_task_metrics_batch_size_avg` and `kafka_connect_connector_task_metrics_batch_size_max` correctly.

Apache patch information

None.

Fixed Issues in Kerberos

There are no fixed issues for kerberos in Cloudera Runtime 7.1.9 SP1.

Fixed Issues in Key HSM

Review the list of Key HSM issues that are resolved in Cloudera Runtime 7.1.9 SP1.

KT-7540: Keytrustee-keyhsm - Upgrade protobuf-java to 3.16.3/3.19.6/3.20.3/3.21.7 due to CVE-2022-3171

Upgraded protobuf-java version to 3.16.3/3.19.6/3.20.3/3.21.7 for keyhsm due to CVE.

KT-7536: Keytrustee-keyhsm - Upgrade Jetty to 9.4.54.v20240208 due to CVE-2024-22201

Upgraded Jetty version to 9.4.54.v20240208 due to CVE

KT-7530: Keytrustee-keyhsm - Upgrade Spring Framework to 6.1.6/6.0.19/5.3.34 due to CVE-2024-22243, CVE-2024-22259 and CVE-2024-22262

Upgraded Spring Framework version to 6.1.6/6.0.19/5.3.34 for keyhsm due to CVE

Fixed Issues in Key Trustee KMS

Review the list of Key Trustee KMS issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-69936: KEYTRUSTEE_SERVER-1 first run failure on 7.1.9 SP1 ("Composite command feature flag is not enabled")

Fixed as part of OPSAPS-70560

Fixed Issues in Kudu

Review the list of Kudu issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-67914 : Kudu - 7.1.9 SP1 Supportability Backports

Fixed Kudu supportability issues.

CDPD-67911 : Critical Bug Fix Backports

Fixed critical issues.

Apache patch information

Supportability fixes in Kudu:

- KUDU-3402: Fixed tracing.html to work with the newer versions of browsers.
- KUDU-3455: Optimized C++ Kudu client when pruning hash partition for in-list predicate.
- KUDU-3487:
 - No more leader step down is performed during rebalancing if the replication factor is 1.
 - Enhancements are made to Kudu diagnostic parse_metrics to differentiate between the values not present and zero value.
 - Added a new metric acceptor_dispatch_times to track dispatching times of newly accepted connections by AcceptorPool.
 - Optimized verification of security related parameters.
 - Expose the most recently seen value in histogram for better troubleshooting.
 - Increased the value of --rpc_acceptor_listen_backlog to 512 to help busy Kudu clusters with larger bursts of incoming requests.
 - Introduced metrics --rpc_listened_socket_rx_queue_size and --rpc_pending_connections to understand and debug issues at network level.
 - Fixed a rare condition that occurred during --rpc_pending_connections metric collection. Optimized memory usage by not allocating memory for undo deltas.
- KUDU-3337: The create_unsafe_meta parameter is now added to aid during power outage situations, where metadata of a tablet is lost. Slow DNS threshold logging for Java Kudu client is now 200 ms. Added additional HTTP headers into Kudu embedded webserver's responses to avoid getting flagged by security scanners.
- KUDU-3534: Additional logging is not performed to log timestamp of two matching DELETE REDO mutations. Added sequenceId into Remote Procedure Call (RPC) trace of Java Kudu client to provide better tracking of retries performed.

- KUDU-3518: Fixed the addition of a new unexpected scan predicate on a primary key column when scanning Kudu tables.
- KUDU-3561: As expected tablet entities in Prometheus format are now skipped without any warning.
- KUDU-3566: Corrected the output of various Kudu metrics backed by High Dynamic Range (HDR) histograms. Added a new option kudu fs in the tool, to locate a given block on file system by the block's ID.
- KUDU-3569: Fixed race condition when optimizing primary key predicates.
- KUDU-3570: Addressed heap-use-after-free and data race issues reported in the MajorDeltaCompactionOp operation.

Critical fixes in Kudu:

- KUDU-3448: It is now possible to encrypt IPKI and TSK private keys.
- KUDU-3500: Optimized Kudu writes by not starting write operations expired in the prepare queue.
- KUDU-3495: Fixed issue where writing using stale Kudu table object resulted in unexpected results.
- KUDU-3461: Avoid Impala crash by returning error if invalid tablet ID is found.
- KUDU-3516: Optimized the process of picking next delta store to be compacted.
- KUDU-3514: Added per-RPC metric to track timeouts.
- KUDU-3483: Fixed flushing data in AUTO_FLUSH_BACKGROUND mode when table schema is stale.
- KUDU-3521:
 - Fixed server crash when the clock was synchronized by PTPd.
 - Improved logging on rejected write requests to help troubleshooting.
 - In the Kudu CLI tool, the `local_replica copy_from_remote` command now returns a non-OK status when tablet copy operations fail.
- KUDU-3535: Fixed a concurrent problem while traversing the traces of Kudu RPC in the Java Kudu client
- KUDU-3526: Scanner in Java Kudu client now binds with a tserver. Fixed NullPointerException (NPE) in Java Kudu client that occurred in busy clusters.
- KUDU-3558: The tool now displays errors early when keytab file flag is empty without starting the Ranger subprocess.
- KUDU-3549: Fixed Prometheus metrics for non-arithmetic gauges.
- KUDU-3486: Tombstoned replica reports are now sent periodically to delete other replicas.
- KUDU-3562: Fixed integer overflow in available space metrics.
- KUDU-3564: Fixed issue where table with range specific hash schema when queried with InList predicate lead to incorrect results.

Fixed Issues in Knox

Review the list of Knox issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-68193: Exclude services/roles from being discovered

New gateway-site.xml properties were added to exclude certain services and roles from being discovered during CM service discovery. The property names are `gateway.cloudera.manager.service.discovery.excluded.service.types` and `gateway.cloudera.manager.service.discovery.excluded.role.types`.

CDPD-68084: [7.1.x] KnoxCLI command for generating descriptor for a role type from a list of hosts

Adding a command to `knoxccli` that can generate a topology descriptor from a list of URLs. The command usage/parameters are described in <https://github.com/apache/knox/pull/835>.

CDPD-68065: [7.1.x] - Add configurable socket / read timeout parameter to discovery client

The following gateway-site.xml properties were added for better control over CM service discovery timeouts in Knox:

- `gateway.cloudera.manager.service.discovery.connect.timeout.ms`
- `gateway.cloudera.manager.service.discovery.connect.read.ms`
- `gateway.cloudera.manager.service.discovery.connect.write.ms`

CDPD-67816: [Analyze] [Knox] [ST] Multiple tests fail on fips due one of the Knox instance restart fail

The Knox Gateway may not start due to a `NullPointerException` in the `ClouderaManagerClusterConfigurationMonitor` class. This could happen when the previously persisted CM cluster config file is empty. This ticket addresses that issue. There is an easy workaround: remove the content of the `$KNOX_DATA_DIR/cm_clusters` folder by running the `'rm -f /var/lib/knox/gateway/data/cm-clusters/*'` command.

CDPD-67600: Knox - Upgrade postgresql to 42.5.5/42.6.1/42.7.2 due to CVE-2024-1597

Upgraded postgresql to 42.5.5/42.6.1/42.7.2 due to CVE-2024-1597.

CDPD-67568: Intermittent HTTP 401 error codes in Oozie tests due to Shiro unable to login: null

Switched back to `org.kushuke.libpam4j` dependency because the one from `org.glassfish.main.libpam4j` produced unforeseen errors.

CDPD-65490: The Knox topology did not update within 600 seconds

There was an issue while Knox compared the generated XML topology with the one that is currently deployed, in highly concurrent environments. The fix was to make the code, that parses the new topology, `threadsafe`.

CDPD-60630: Knox was redirecting YARN Node Manager URLs to HTTP instead of HTTPS

The YARN service definition was modified such that it no longer hard-codes the HTTP scheme in any of its URL rewrite rules, making it applicable to proxying YARN deployed with and without TLS.

Apache patch information

- KNOX-3002
- KNOX-3026
- KNOX-3000
- KNOX-3022

Fixed Issues in Livy

There are no fixed issues for Livy in Cloudera Runtime 7.1.9 SP1.

Fixed Issues in MapReduce

There are no fixed issues for MapReduces in Cloudera Runtime 7.1.9 SP1.

Fixed Issues in Navigator Encrypt

Review the list of Navigator Encrypt issues that are resolved in Cloudera Runtime 7.1.9 SP1.

KT-7503: Navigator Encrypt, when using KMS, does not do kinit on host reboot

When working with Ranger KMS, it is a requirement that user root obtain a Kerberos ticket otherwise any Navigator Encrypt request to Ranger KMS will fail with an `Unauthorized request` error. If a ticket does not exist during a host reboot, `NavEncrypt` will be unable to mount any drives. This means when a host reboots `NavEncrypt` needs to obtain a Kerberos ticket.

KT-7513: Navigator Encrypt needs to generate a Kerberos keytab file when working with Ranger KMS

There is a new script that builds a Kerberos keytab file if you are using Navigator Encrypt with Ranger KMS. It is called `navencrypt-gen-keytab` and invokes Cloudera Manager to create the keytab.

This script works with Cloudera Manager to create a keytab file to be used by Navigator Encrypt. It accepts the URL of Cloudera Manager and credentials that works with Cloudera Manager. The script format is as follows:

```
navencrypt-gen-keytab <cm-url> <user:pw>
```

KT-7519: When removing the kernel module, Navigator Encrypt needs better version checking

When removing the kernel module from the system, Navigator Encrypt now makes better version checks.

KT-7528: Log message emitted when drive is mounted with Navigator Encrypt

When a drive is mounted by Navigator Encrypt (by using navencrypt-prepare) the log message emitted is useless. This is what it looks like:

```
2024-02-22 17:47:32,342 +0000 level=INFO app=mount-navencrypt action=mount 11
```

This issue is fixed now.

KT-7529: Logging issues in Navigator Encrypt

Currently, the message "Registration was successful" is emitted too soon before all the operations for a successful registration are completed. Now this message is emitted after all deposits are written to the Key Manager.

Currently all messages in the navencrypt log include a timestamp with "+0000" in the timestamp, which indicates GMT. Now the log entries look like the following format with the new timestamp:

```
2024-02-28T11:53:01-0800 level=INFO app=navencrypt action=control  
Fetching deposit from Key Manager  
2024-02-28T11:53:06-0800 level=INFO app=navencrypt action=control  
Deposit fetched successfully
```

KT-7559: Page allocation failures in kernel memory

A problem has been fixed whereby Navigator Encrypt was causing the kernel to hang while obtaining kernel memory.

Due to this fix, the command "navencrypt acl --update" has to be run before the 7.1.9.1000 version is started.

ENGESC-25935: Page allocation failure

If a customer is running RHEL 8.8 (which is Linux Kernel 4.18), it is possible for the kernel to report a page allocation failure and cause a kernel panic.

This issue is fixed now.

Fixed Issues in Oozie

Review the list of Oozie issues that are resolved in Cloudera Runtime 7.1.9 SP1.

Apache patch information

None

Fixed issues in Ozone

Review the list of Ozone issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-68972: When a leader Ozone Manager election gets stuck, the leader OM is not available. This causes the failure of the OM operation.

This issue is now resolved.

CDPD-61426: Improve Ozone metadata performance by avoiding duplicated internal lookups.

This issue is now resolved.

CDPD-57118: Improve Ozone metadata performance by optimizing internal locking implementation.

This issue is now resolved.

CDPD-69167: Upgrade Bouncy Castle to 1.78.

Upgraded Bouncy Castle Java release version to 1.78 to address issues identified in CVE-2023-33202, CVE-2023-33201, CVE-2024-29857, CVE-2024-30171, and CVE-2024-30172.

CDPD-68255: Upgrade Commons-configuration2 to 2.10.1.

Upgrade Commons-configuration2 version to 2.10.1 to address issues identified in CVE-2024-29133 and CVE-2024-29131.

CDPD-68159: File encryption information is not properly set when overwriting a file.

When a file is overwritten, the new version of the file uses encryption information of the old version file. This issue is now resolved.

CDPD-62273: Incorrect error message is displayed when deleting a file in a .snapshot path.

Correct error message is displayed now when deleting a file in the .snapshot path.

CDPD-62272: Return error response when creating a directory in the .snapshot path.

This issue has been fixed by displaying an error if the directories in the the reserved .snapshot path are deleted.

CDPD-61779: SnapshotDiff cannot accept continuation token greater than the length of the actual snapdiff.

This issue has been fixed by displaying an invalid token exception if the token is out of the range.

CDPD-67224: Upgrade Spring Framework to 6.1.6/6.0.19/5.3.34.

Upgrading spring framework to 5.3.34 to address issues identified in CVE-2024-22243, CVE-2024-22259, and CVE-2024-22262.

CDPD-67146: Add permission check for snapshot diff.

Only users with bucket read permission are allowed to get snapshot diff between two snapshots.

CDPD-64545: Recon Backport HDDS-9599 to CDP Private Cloud Base 7.1.9 SP1.

This issue is now resolved by highlighting the search icon in the OM DB Insights of Recon to navigate to containers.

CDPD-64540: Recon Backport HDDS-9161 to CDP Private Cloud Base 7.1.9 SP1.

This issue is now resolved.

CDPD-64500: Option to enable Ratis in SCM.

Starting CDP Private Cloud Base 7.1.9 SP1, you can enable Ratis in SCM on an existing Non-HA cluster by modifying ozone.scm.ratis.enable property. Enabling Ratis is prerequisite for making SCM Highly Available. This issue is now resolved.

CDPD-64285: Implement snapshot list pagination.

This issue is now resolved by adding the listOptions to list snapshot CLI to make the pagination work correctly.

CDPD-63165: The InvalidStateTransitionException exception is displayed after recommissioning SCM.

When a new SCM is bootstrapped with Container CLOSE transactions in Raft Log for the containers that are closed due to Pipeline failure, the newly bootstrapped SCM will fail to start as those containers are in OPEN state on the new SCM. To avoid this, the finalization of containers is now done only by the leader SCM and this operation goes through Ratis so that the followers also update the container state in the DB. This issue is now resolved.

CDPD-56768: ReplicationFactor may be one or three.

This issue occurs when distcp tries to create target replicas with a replication factor of 1 when the source replicas were Erasure Coded. This issue is now resolved.

Technical Service Bulletins

TSB 2024-749: Possible Ozone Snapshot Chain Corruption in 7.1.9.0

For the latest update on this issue see the corresponding Knowledge article: [TSB 2024-749: Possible Ozone Snapshot Chain Corruption in 7.1.9.0](#)

Apache Patch Information

- HDDS-10594

- HDDS-9423
- HDDS-9422
- HDDS-9360
- HDDS-10434
- HDDS-9599
- HDDS-9161
- HDDS-9897
- HDDS-9832
- HDDS-10231
- HDFS-17376
- HDDS-8765
- HDDS-9297

Fixed Issues in ORC

Review the list of ORC issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-69642: Use Centralised Guava version in ORC

Upgraded guava library from version 28.1-jre to 32.0.1-jre due to CVE-2023-2976

Fixed Issues in Parquet

There are no fixed issues for Parquet in Cloudera Runtime 7.1.9 SP1.

Fixed Issues in Phoenix

Review the list of Phoenix issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-69580: Backport of PHOENIX-6687

This is a backport of PHOENIX-6687 - the region server hosting the SYSTEM.CATALOG file fails to serve any metadata requests as the default handler pool threads are exhausted.

CDPD-60862: DDL refers to "Data Definition Language", a subset of SQL statements that change the structure of the database schema in some way, typically by creating, deleting, or modifying schema objects such as databases, tables, and views. Cloudera services that support DDL include Impala, Hive (using HiveQL), Spark (using SparkSQL), HBase, Phoenix, Flink, and Kafka.

This issue is fixed.

CDPD-35925: Omid service fails to start rarely. Restarting it solves the problem most of the times.

This issue is fixed.

Apache Patch Information

- PHOENIX-6687

Fixed Issues in Ranger

Review the list of Ranger issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-70349: Ranger admin's service/metrics/status returns 404 and throws error in catalina.out

NA

CDPD-70423: Ranger replication failure during ranger-replication-transform

Ranger replication policy fails in the "Transforming services, policies and roles from Ranger" step during the policy execution. The fix has been provided for this scenario.

CDPD-68920: Custom Service Descriptors (CSD) changes for RMS metrics for Hive and Ozone service mappings

Ranger RMS metrics was previously supported only for Hive and HDFS service mappings. Ranger RMS now supports multiple source service types. Ranger RMS provides mappings for Hive-HDFS

as well as Hive-Ozone services, where Hive, HDFS, Ozone, Ranger, and Ranger RMS services are installed and RMS is enabled.

CDPD-68831: Ranger react UI - Some modules showed hard coded time zone string like "Indian Standard Time"

Ranger react UI popup was showing details of an admin audit log that had a hard coded string namely "Indian Standard Time". Certain logs table headers also had hard coded string namely "Indian Standard Time". The fix has been provided for this scenario.

CDPD-68807: Upgrade Spring Framework to 6.1.6/6.0.19/5.3.34 due to CVE-2024-22243, CVE-2024-22259 and CVE-2024-22262

Upgrade Spring Framework to 5.3.34

CDPD-68737: Upgrade opensearch to 1.3.15 due to CVE-2023-45807

Upgrade opensearch to 1.3.15

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

In Ranger replication Resource and Hive URL mapping, there are different paths:

For example:

/user/hive/warehouse/

s3a://bucket/path

Including a / separator or not can cause issues during transformation.

Fix is provided for Resource and Hive URL prefixed based transformations to be aware of potential / signs at the end of URLs and handle the transformation correctly.

CDPD-68590: 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 who is the owner of the Hive database for the attempted operation is allowed.

CDPD-68489: Upgrade jline to 3.25.1 due to CVE-2023-50572

Upgraded jline to 3.25.1 due to CVE-2023-50572

CDPD-68375: Enhance handling of subAccess authorization in Ranger HDFS plugin

Currently, Ranger performs authorization of the HDFS commands which require access to the hierarchy of files/directory rooted at the argument passed to the HDFS command.

CDPD-68365: Audit logs for tag masking policy is displaying the service name along with datamask type

A fix is provided to not display the service type name along with datamask type in audit logs for tag based masking policy.

CDPD-68252: Upgrade Commons-configuration2 to 2.10.1 due to CVE-2024-29133 and CVE-2024-29131

Upgraded Commons-configuration to v2.10.1 due to CVEs

CDPD-68024: Use name validation regex instead of service name validation regex for display name

Currently, in the Service creation / edit form in Ranger react UI, there is a validation for the display name which does not allow spaces. The validation has to be modified to include spaces, otherwise it does allow you to edit the Service from the UI.

CDPD-67864: Upgrade Spring Security to 5.7.12/5.8.11/6.1.8/6.2.3 due to CVE-2024-22257

Upgrade Spring security to 5.7.12

CDPD-67750: Upgrade telemetry to 1.36.0

Upgrade telemetry version to 1.36.0

CDPD-67749: Upgrade protobuf-java to 3.21.7 due to CVE-2022-3171

Upgrade protobuf-java to 3.21.7

CDPD-67746: Upgrade Nimbus-JOSE-JWT to 9.37.3 due to CVE-2023-52428

Upgrade Nimbus-JOSE-JWT to 9.31

CDPD-67744: Exclude Apache Derby from ranger-rms module due to CVE-2022-46337

Exclude Apache Derby from the ranger-rms module

CDPD-67530: Support FIPS on JDK17

Successfully validated Ranger and Ranger KMS for FIPS with JDK 17.

CDPD-67316: Add an eye icon for password visibility in Ranger login page

The Ranger login page did not support password visibility functionality. Users were not able to verify the entered password. Users can now view the entered password by clicking on the eye icon on the password field in the Ranger login page.

CDPD-67315: Filtering the resources in the search filter options on the policy listing page based on policy type.

N/A

CDPD-67314: Handling local storage data for column show/hide functionality

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

CDPD-67110: Upgrade json-smart due to CVE-2023-1370

Upgraded json-smart to 2.4.10 due to CVEs

CDPD-67098: Upgrade commons-compress to 1.26.0 due to CVE-2024-25710 and CVE-2024-26308

Upgraded commons-compress to 1.26.0 due to CVE-2024-25710 and CVE-2024-26308

CDPD-66512: Enabling FIPS for Luna 7 server

A fix was provided.

CDPD-66501: Remove/Replace OpenSAML v3 due to EOL

Removed OpenSAML v3.4.5 due to EOL

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

The Ranger Admin UI's "Plugin Status" page now supports pagination. Previously, the Ranger Admin UI displayed only (the first or random) 200 entries, which was not easy to navigate. Also, the ordering/sorting on columns was only a "client side" sorting which hampered viewing some entries (only if searched by Host Name for example, it was possible). With bigger clusters it was a usability issue.

CDPD-65125: User Agent info not logged under "Login sessions" when login fails

User agent information will not be available for failed login.

CDPD-64634: Update ServiceDef icons on the service manager page

The service manager page previously used a folder icon for displaying ServiceDefs. Presently, the service manager page displays a real logo or icon that is specific to each ServiceDef.

CDPD-64334: Ranger - Upgrade Bouncy Castle to 1.77 due to CVE-2023-33202 and CVE-2023-33201

Bouncy Castle lib version upgraded to v1.77

CDPD-63694: Ranger plugin sends audits with Public Key Infrastructure (PKI) authentication with Solr user and causing PKI Exception

The error occurs when the workload Solr sends the audits to the infra Solr. The Ranger plugins' own solrclient uses PKI whatever the Solr configuration is.

CDPD-63984: HBase shell revoke command failed with 'HTTP 400 Error: processSecureRevokeRequest processing failed'

A fix was provided for this issue.

CDPD-62362: Ranger plugin sends audit information with Public Key Infrastructure (PKI) authentication with Solr user and causing PKI Exception

A fix is provided to handle the error when the workload Solr sends the audit information to the infra Solr.

CDPD-65125: User Agent information not logged under "Login sessions" when login fails

A fix is provided where the User Agent information is available for failed login.

CDPD-61934: Filter audits for cc_metric_reporter user on Kafka service repository

A lot of audits are generated for cc_metric_reporter user for the Kafka service repository for the resource "__CruiseControlMetrics". These audits will fill up the audit logs, and not much value is added by these audits. A fix is provided by adding Filter audits for cc_metric_reporter user on Kafka service repository.

Apache patch information

- RANGER-4504
- RANGER-4729
- RANGER-4780
- RANGER-4753
- RANGER-4673
- RANGER-4519
- RANGER-4348
- RANGER-3346
- RANGER-4467
- RANGER-4659
- RANGER-4656
- RANGER-4641

Fixed Issues in Ranger KMS

Review the list of Ranger issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-58035: Luna HSM w/ Hardware in FIPS mode fails

When using a Luna HSM with firmware version of 7.7.2 or newer in FIPS mode, the HSM fails to encrypt.

This issue is fixed now.

Fixed Issues in Schema Registry

There are no fixed issues for Streams Replication Manager in Cloudera Runtime 7.1.9 SP1.

Fixed Issues in Solr

There are no fixed issues for Solr in Cloudera Runtime 7.1.9 SP1.

Apache patch information

None

Fixed Issues in Spark

There are no fixed issues for Spark in Cloudera Runtime 7.1.9 SP1.

Apache Patch Information**Fixed Issues in Spark Atlas Connector**

Review the list of Spark Atlas Connector issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-53836: Request for fault tolerant Spark atlas hook

Added fault tolerant mode to the Spark Atlas connectors. For more information see the [Spark connector configuration in Apache Atlas](#) documentation.

Apache Patch Information

Fixed Issues in Sqoop

There are no fixed issues for Sqoop in Cloudera Runtime 7.1.9 SP1.

CDPD-59557: Secure options to provide the Hive password for Sqoop Hive imports

This fix introduces secure options that you can use to provide the Hive password during Sqoop-Hive imports instead of the earlier way of providing the password as plaintext in the command-line interface.

Fixed Issues in Streams Replication Manager

There are no fixed issues for Streams Replication Manager in Cloudera Runtime 7.1.9 SP1.

Fixed Issues in Streams Messaging Manager

There are no fixed issues for SMM in Cloudera Runtime 7.1.9 SP1.

Fixed Issues in Tez

There are no fixed issues for Tez in Cloudera Runtime 7.1.9 SP1.

Apache Patch Information

None

Fixed Issues in YARN and YARN Queue Manager

Review the list of YARN and YARN Queue Manager issues that are resolved in Cloudera Runtime 7.1.9 SP1.

COMPX-16696: YARN Queue Manager - Dropwizard version upgrade to 2.1.11

Dropwizard is upgraded to version 2.1.11 from an older, vulnerable version. After upgrade, YARN Queue Manager now does not disable the replacement rules.

COMPX-16626: YARN Queue Manager - Upgrade Bouncy Castle version to 1.78 due to CVE-2024-29857, CVE-2024-30171 and CVE-2024-30172

Upgrade the Bouncy Castle version to 1.78 to fix the following vulnerabilities.

- CVE-2024-29857: Importing an EC certificate with specially crafted F2m parameters can cause high CPU usage during parameter evaluation.
- CVE-2024-30171: Possible timing-based leakage in RSA-based handshakes due to exception processing eliminated.
- CVE-2024-30172: A crafted signature and public key can be used to trigger an infinite loop in the Ed25519 verification code.

COMPX-16386: CDPD - Upgrade Spring Security version to 5.7.12/5.8.11/6.1.8/6.2.3 due to CVE-2023-34042 and CVE-2024-22257

Upgrade the Spring Security version to 5.7.12/5.8.11/6.1.8/6.2.3 to fix the following vulnerabilities.

- CVE-2023-34042: The `spring-security.xsd` file inside the `spring-security-config` JAR file is world-writable. If extracted, it could be written by anyone with access to the file system. While there are no known exploits currently, this is an example of “CWE-732: Incorrect Permission Assignment for Critical Resource” and could potentially result in an exploit. You should update to the latest version of Spring Security to mitigate any future risks associated with this issue.
- CVE-2024-22257: In Spring Security, versions 5.7.x prior to 5.7.12, 5.8.x prior to 5.8.11, 6.0.x prior to 6.0.9, 6.1.x prior to 6.1.8, and 6.2.x prior to 6.2.3, an application is potentially vulnerable to broken access control when it directly uses the `AuthenticatedVoter#vote` method with a null `Authentication` parameter.

COMPX-11626: Improved management of HDFS delegation tokens in YARN

YARN now efficiently manages HDFS delegation tokens before log aggregation attempts to write logs to HDFS. Tokens are iterated through and renewed as necessary. If renewal fails and a token expires, it is promptly removed from the credentials, ensuring smoother operation of log aggregation processes.

COMPX-11263: Inconsistent display of updated configuration values in YARN Queue Manager UI

The YARN Queue Manager UI intermittently displayed old configuration values after updates. This issue has been fixed, and the UI now reliably shows only the new values entered by the user immediately after an update.

CDPD-66169: Apache Hadoop - Upgrade of bcpkix-jdk15on library version to address CVE-2019-17359

The bcpkix-jdk15on library has been upgraded to version 1.70 in Hadoop to address the security vulnerability identified in CVE-2019-17359.

CDPD-60630: Knox redirecting YARN Node Manager URLs to HTTP instead of HTTPS

The YARN service definition is modified to remove the hard-coded HTTP scheme from its URL rewrite rules. This change ensures compatibility with proxying YARN whether it is deployed with or without TLS, resolving the issue of Knox redirecting YARN Node Manager URLs to HTTP instead of HTTPS.

Apache patch information

None

Fixed Issues in Zeppelin

Review the list of Zeppelin issues that are resolved in Cloudera Runtime 7.1.9 SP1.

CDPD-68796: Zeppelin - Upgrade Apache Maven to 3.8.6 due to CVE-2021-26291

Upgraded to Apache Maven 3.8.6 to resolve CVE-2021-26291 by no longer following HTTP (non-SSL) repository references in POM files, thereby mitigating the risks of malicious code injection.

Fixed Issues in Zookeeper

Review the list of Zookeeper issues that are resolved in Cloudera Runtime 7.1.9 SP1.

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

Apache ZooKeeper has fixed an information disclosure vulnerability (CVE-2024-23944) caused by a missing ACL check in persistent watcher handling. Attackers may exploit this by attaching a persistent watcher to a parent znode, allowing them to monitor child znode's without ACL verification, exposing the full path of znodes. This path may contain sensitive information, like usernames or login IDs. The issue is fixed in ZOOKEEPER-4799.

Apache patch information

- ZOOKEEPER-4799

Known issues in Cloudera Runtime 7.1.9 SP1

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.1.9 SP1.



Note: CDSW does not support RPM-based installation on CDP Private Base. (RPM installation is deprecated and only supported on HDP and CDH 5. For CDH6 and onward, Cloudera recommends you to use CSD-based installations.)

Known Issues in Atlas

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

CDPD-6565: Issue with ddlQueries in Atlas for Hive/Impala tables created in default/non-default db using DAS/HUE respectively

ddlQueries are not created for Impala origin table. ddlQueries created for Impala/Hive CTAS tables have difference in names as compared to same operation for Hive table. The Impala table has the query in the ddlQueries entity name.

None

CDPD-55301: ddlQueries and ALTERNATIVE_* lineage missing for Spark tables created through spark3-shell

The ddlQueries and ALTERNATIVE_* lineage missing for Spark tables created using spark3-shell.

None

CDPD-56085: 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

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*.

None

CDPD-58581: storage_handler is not set in Atlas for Impala to Iceberg in-place migrated tables

The storage_handler property is not set for Iceberg tables created in Impala, because in-place migration is not supported in Impala for the current release.

None

CDPD-58554: Discard audits of specific classification, label, and business metadata

Support to control audits for specific classification, label, and business metadata is not present in the Custom Audit Filters feature.

None

OPSAPS-67783: During rolling upgrade one among two Atlas server failed to start but Cloudera Manager considered as success

Cloudera Manager marks the Execute command Start on service Atlas-1 as a success even when Atlas service had failed to start successfully. In such cases, Atlas logs give the exact reason for Atlas start-up failure.

None

CDPD-19358: "IsIndexable" and "isOptional" value of a typedef's attribute is modified post migration.

None.

CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command

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

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

In converting content from Navigator to Atlas, the create time for Hive tables is not moved to Atlas.

None

CDPD-10574: Suggestion order doesn't match search weights

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

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) does not have the clusterName appended to it as do hdfs_path entities in other location types.

None.

CDPD-5933, CDPD-5931: Unexpected search results when using regular expressions in basic searches on classifications

When you include a regular expression or wildcard in the search criteria for a classification in Basic Search, the results may differ unexpectedly from when full classification names are included. For example, the Exclude sub-classifications option is respected when using a full classification name as the search criteria; when using part of the classification name and the wildcard (*) with Exclude sub-classifications turned off, entities marked with sub-classifications are not included in the results. Other instances of unexpected results include case-sensitivity.

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 the metadata collection from HMS. For example, if an ALTER TABLE operation in Spark that is changing a table name and is reporting 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 "@" does not fetch the correct results

When searching Atlas qualifiedName values that include an "at" 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, and use 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 does 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 logout 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 CDP.

None.

CDPD-1664: Guest users are redirected incorrectly

Authenticated users logging in to Atlas are redirected to the CDP 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: The 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, to add metadata tags to relationships that you wanted to make sure were unique in the system, you can design the relationship attribute with the property IsUnique equal true. However, in this release, the IsUnique attribute is not enforced.

None.

DOCS-13759: Tag Propagation stops after a certain depth while the lineage is being extended

When a tag is added to an entity at timestamp T1, the entities along the lineage to which the tag must be propagated is calculated at T1. Before tag propagation completes, if the lineage is extended, tag does not propagate to the entities in the extended lineage.

DOCS-13760: System Attributes search, __classificationNames: Search with parent tag does not return entities associated to its children tags

System attribute search with __classificationNames = parent_tag returns entities associated to parent_tag only and not entities associated to its children tag.

Instead of using system attribute, employ the basic search attribute classification which lists entities associated with inherited classifications.

CDPD-41142: When a Kafka console consumer group is run, more than one update audits are seen

After running the console consumer with a consumer group, verify the consumer group entity created, along with the metrics and notifications for the consumer group and topic. The expected result can be: one ENTITY_CREATE audit and one ENTITY_UPDATE audit. But more than one ENTITY_UPDATE audits are seen.

CDPD-40165: Two audits are created for SPARK CTAS table

When following Spark queries are fired:

```
spark.sql("create table table1(id int)")
spark.sql("create table table2 as select * from table1")
```

HMS sends "ENTITY_CREATE" and "ENTITY_FULL_UPDATE_V2".

The extra ENTITY_FULL_UPDATE_V2 message received from HMS is sent as a part of ALTER_TABLE_ADDCOLS event from the HMS Hook side. This behaviour is observed only when the queries are run from Spark SQL and not when run the same queries are run from Beeline.

CDPD-39197: Debug metrics returns empty data

When debug metrics is enabled and some operations are performed, the response is empty

CDPD-36495: Updating legacyAttribute from False to True resets the initially created relationshipAttributes values

Creating types, entities, and to start, you must set the relationship with is_legacy_attribute value as **False**.

Later, update the value relationshipDef is_legacy_attribute to True.

For the entities that were created before updating the is_legacy_attribute to True, the relationship Attributes value is reset.

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-29409: Hive import: Suggestion suggests entity which is deleted.

Suggestions suggests tables of a database, which is a deleted entity.

CDPD-25152: Tag propagation through deferred actions consumes additional time as compared to default flow

The additional time might be due to the small overhead added to create / update task vertex and which is run in the background. This also depends on number of tasks queued to be executed.

CDPD-42954: Zeppelin notebook fails after enabling Atlas-HDFS hook

The Zeppelin notebooks are failing with errors after enabling Atlas-HDFS hook in the CDP cluster.

When the below properties are set for atlas-client.properties in Cloudera Manager:

- atlas.jaas.KafkaClient.option.keyTab
- atlas.jaas.KafkaClient.option.principal

Along with adding the properties in /etc/atlas/conf/atlas-application.properties, Cloudera Manager also adds these properties to atlas-application.properties for other services (like Spark).

Adding these properties interferes with the normal flow of the services (like Spark)

To enable HDFS lineage feature, instead of setting these properties through Cloudera Manager, users can manually add the properties directly in /etc/atlas/conf/atlas-application.properties

CDPD-40346: The ddlQueries and ALTERTABLE_ADDCOLS lineage missing for Impala tables.

The ALTERTABLE_ADDCOLS lineage has some issue when an Impala table is altered and the corresponding lineage is not created.

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

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

None

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

Adding Japanese labels or user defined properties results in the error message: “Invalid label: ###, label should contain alphanumeric characters, _ or -”

None

CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties

Entities, such as `hdfs_path`, are not returned by the search when their indexable properties are searched using partial search terms made of Japanese or Chinese characters. Only exact matches return results when searching the indexed properties made of Japanese or Chinese characters.

None

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

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

None.

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

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

None.

CDPD-71482: BCFKS keystore type is not recognized by the Atlas client

The Schema Registry - Atlas integration fail on the FIPS environment because the BCFKS keystore type is not recognized by the Atlas client. Currently, it is hard-coded as "JKS".

None.

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

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

None.

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

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

None.

CDPD-43772: Performance issues with Atlas service

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

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

CDPD-74467: Ranger TagSync does not support ofs path parsing for Ozone-Atlas

When trying to create Hive table on top of Ozone ofs path with the command `ofs://ozone1/volume1/bucket1/table_oz_demo1`, Ranger TagSync fails with an error.

Upgrade to Cloudera Private Cloud Base 7.1.9 SP1 CHF3.

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet `atlas.authentication.method.file` is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

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

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-77911: Missing Log4j Redactor dependency

The missing `org.cloudera.logredactor` dependency can result impossible Log4j errors during startup. Additionally, it can lead to unwanted data leak because of incorrect log redaction.

CDPD-79160: NPE while deleting BusinessMetadata

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

None

Apache Jira: [ATLAS-4863](#)

CDPD-80921: Without permission for one glossary, /glossary call throws exception

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-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export

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

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](#)

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

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

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error `"NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl"`.

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

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

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

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

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

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-88230: Bulk import operations currently do not construct lineage as expected

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

Bulk import operations do not create audits for the imported entities.

Known Issues in Avro

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

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

None.

Known issues in Cloud Connectors

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

CDPD-76378 Uploading files to S3 takes longer than expected

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.

Rebalancing with Cruise Control does not work due to the metric reporter failing to report the CPU usage metric

On the Kafka broker, the Cruise control metric reporter plugin may fail to report the CPU usage metric.

If the CPU usage metric is not reported, the `numValidWindows` in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.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 the [Moving a Host Between Clusters](#).



Note: Cluster nodes affected by this issue are not displayed as unhealthy.

OPSAPS-69978: Cruise Control capacity.py script fails on Python 3

Cruise Control might fail to start on Python 3 when capacity information is queried during the startup process. This is caused by a breaking change between Python 2 and 3.

None

Known Issues in Hadoop

There are no known issues for Hadoop in Cloudera Runtime 7.1.9 SP1.

Known Issues in HBase

This topic describes known issues and workarounds for using HBase in Cloudera Runtime 7.1.9 SP1.

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

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

The following services support DDL statements:

- Impala
- Hive – using HiveQL

- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

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

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

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

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

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

HDFS encryption with HBase

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

Workaround: N/A

AccessController postOperation problems in asynchronous operations

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

- If a Delete Table fails for a reason other than missing permissions, the access rights are removed but the table may still exist and may be used again.
- If `hbaseAdmin.modifyTable()` is used to delete column families, the rights are not removed from the Access Control List (ACL) table. The `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.

Workaround: N/A

Apache Issue: [HBASE-6992](#)

Snappy compression with /tmp directory mounted with noexec option

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

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

HBase shutdown can lead to inconsistencies in META

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

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

```
hbase master stop --shutdownCluster
```

The command output must end with Closing master protocol: MasterService phrase. You can verify the command execution by checking the master logs. The log must contain Cluster shutdown req

uested 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.

The custom JWT Principal claim is not usable for HBase

The `hbase.security.oauth.jwt.token.principal.claim` configuration property does not exist in the 7.1.9 SP1 version. The purpose of this property is to allow the use of a Subject or Principal claim different from the default sub claim.

Workaround: Use the standard sub claim in the JWT token.

HBase procedures might hang because WriteAheadLog for ProcedureStore does not progress

HBase Master might hang and cause longer recovery when the WriteAheadLog is stuck while persisting to the ProcedureStore. HBase Master eventually aborts itself, however, the operations that rely on the ProcedureStore (for example, region movements and table creation) do not progress during this period.

Workaround: Restart the active HBase Master to speed up the recovery.

CDPD-77399: HBase fails to register the servlet metrics and throws `ClassNotFoundException: org.apache.hadoop.metrics.MetricsServlet`

The `MetricsServlet` class is a Hadoop 2-based metric servlet unavailable in Hadoop 3 deployments.

Workaround: Ignore this WARN log message during HBase Master and RegionServer startup.

Apache Issue: [HBASE-28315](#)

Known Issues in HDFS

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

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	<code>dfs.permissions.ConflictSummary.subAccess</code>	Not Set	True	Performance optimization for NameNode content summary API
HDFS	<code>dfs.datanode.handler.count</code>	3	10	Optimal value for DN server threads on large clusters

None.

CDPD-60387: Configuration difference between 7.1.8.3 and 7.1.9.0 results

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

None.

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

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

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



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

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

None.

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

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

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

None.

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

None.

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

None.

Unsupported features

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

- ACLs for the NFS gateway ([HADOOP-11004](#))
- Aliyun Cloud Connector ([HADOOP-12756](#))
- Allow HDFS block replicas to be provided by an external storage system ([HDFS-9806](#))
- Consistent standby Serving reads ([HDFS-12943](#))

- Cost-based RPC FairCallQueue ([HDFS-14403](#))
- HDFS Router Based Federation ([HDFS-10467](#))
- NameNode Federation ([HDFS-1052](#))
- NameNode Port-based Selective Encryption ([HDFS-13541](#))
- Non-Volatile Storage Class Memory (SCM) in HDFS Cache Directives ([HDFS-13762](#))
- OpenStack Swift ([HADOOP-8545](#))
- SFTP FileSystem ([HADOOP-5732](#))
- Storage policy satisfier ([HDFS-10285](#))

Known Issues in Hive

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

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

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

The following services support DDL statements:

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

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

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

CDPD-60770: Beeline Authentication Issue with Special Characters in Passwords

When LDAP is enabled, users cannot authenticate with Beeline if the password contains a special character. For example, the following string fails:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>;user=user@XXX;password='R3G#xpXyoy1M0Jb1'
```

Use the -p parameter to execute the Beeline command:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>; -n user@XXX -p 'R3G#xpXyoy1M0Jb1'
```

CDPD-43769: Inconsistent behavior with CHAR comparisons and string literals

In Hive, 'char' data to 'string' data comparison is considered as equal, but 'char' data to 'varchar' data comparison is considered as not equal if the data has trailing spaces. For example,

```
create table test(colchar char(10), colvarchar varchar(20), colstring string);
insert into test values ('a ', 'a ', 'a ');

select count(*) from test where colchar = colstring;
Output returns 1 because char and string data is considered as equal

select count(*) from test where colchar = colvarchar;
```

Output returns 0 because char and varchar data is considered as not equal

It is recommended that you use 'varchar or string' data type instead of 'char' data type for any character or string representation.

CDPD-44060: Issue rebuilding Hive Materialized Views

The ALTER MATERIALIZED VIEW <view_name> REBUILD; command to rebuild Materialized View fails with the following error:

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

This can occur if the Materialized View Rebuild command is abruptly killed or if more than one Materialized View Rebuild command is issued for the same Materialized View. As a result, LOCK entry in the backend metastore DB table 'materialization_rebuild_locks' is not deleted and when the command is issued again, it will fail for that Materialized View.

To resolve this issue, perform the following steps to manually delete the stale locks:

1. From the backend DB, run the following command and check if there are any entries for the view to confirm if stale locks are present:

```
select mrl_db_name, mrl_tbl_name, mrl_last_heartbeat from ma
terialization_rebuild_locks where lower(mrl_tbl_name)='<view
_name>';
```

2. Create a backup of the table if there are multiple entries:

```
create table bkp_top_apps_custom as
select mrl_db_name, mrl_tbl_name, mrl_last_heartbeat from ma
terialization_rebuild_locks where lower(mrl_tbl_name)='<view
_name>';
```

3. Delete the stale lock entries from the table:

```
delete from materialization_rebuild_locks where lower(mrl_tb
l_name)='<view_name>';
```

For more information, see the related [Cloudera Community article](#).

CDPD-42726: Wrong results for agg queries from stats after running multi-insert query on managed table

Multiple inserts at the same time into the same partition results in invalid stats. For example:

```
create table source(p int, key int, value string);
insert into source(p, key, value) values (101,42,'string42');

create table stats_part(key int,value string) partitioned by (p
int);
from source
insert into stats_part select key, value, p
insert into stats_part select key, value, p;

select count(*) from stats_part;
-- In this scenario, StatsOptimizer helps serving this query be
cause the result should be rowNum of the partition p=101.
-- The result is 1, however, the result should be 2.
```

Rewrite the multi-insert statement by merging the branches inserting into the same partition to one branch.

CDPD-57574: Query execution fails due to `NullPointerException` in `DagUtils.setupQuickStart`

Hive queries may fail at runtime with a `NullPointerException` while executing the Tez graph at `DagUtils.setupQuickStart` method. A part of the stack trace from the error is shown below:

```
ERROR : Failed to execute tez graph.
java.lang.NullPointerException: null at org.apache.hadoop.hive.ql.exec.tez.DagUtils.setupQuickStart(DagUtils.java:1724)
```

The error usually occurs from an invalid plan that has cycles created by semi-join or map-join edges.

Disable some of the optimizations in order to remove the cycles from the plan. It is sufficient to set one of the following properties:

- `set hive.tez.dynamic.partition.pruning=false;`
- `set hive.tez.dynamic.semijoin.reduction=false;`
- `set hive.auto.convert.join=false;`

To avoid any undesired performance regressions, disable the optimization(s) only for specific queries and not globally (through Cloudera Manager) for the whole workload or cluster.

CDPD-43107: `SemanticException` for `INSERT INTO` statement when Hive Cost-based Optimizer (CBO) is disabled

If you are running the `INSERT INTO` query and `hive.cbo.enable` is set to "false", the query fails with a `SemanticException`. For example,

```
set hive.cbo.enable=false;

CREATE TABLE mytable (
  id INT,
  str STRING
);

INSERT INTO mytable (id, str) VALUES (1, 'a');
```

```
Output:
org.apache.hadoop.hive.ql.parse.SemanticException: 0:0 Expected
2 columns for insclause-0/default@mytable; select produces 1 col
umns. Error encountered near token 'a'
```

This issue occurs only when columns are specified for the target table in the query. `INSERT INTO mytable values (1, 'a');` does not result in an exception.

Enable CBO and run the query.

CDPD-43957: `HiveServer` shuts down during replication due to high resource usage

During Hive replication (both bootstrap and incremental), you may notice that the `HiveServer` (HS2) shuts down periodically with the following error:

```
java.sql.SQLException: org.apache.hive.jdbc.ZooKeeperHiveClientE
xception: Unable to read HiveServer2 configs from ZooKeeper
at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.ja
va:265)
at org.apache.hive.jdbc.HiveDriver.connect(HiveDriver.java:107)
at java.sql.DriverManager.getConnection(DriverManager.java:664)
at java.sql.DriverManager.getConnection(DriverManager.java:247)
at com.cloudera.enterprise.hive3qt.Hive3QueryTool$HiveOperation.
execute(Hive3QueryTool.java:682)
```

```

    at com.cloudera.enterprise.hive3qt.Hive3QueryTool.main(Hive3QueryTool.java:935)
    Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Unable to read HiveServer2 configs from ZooKeeper
    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:177)
    at org.apache.hive.jdbc.Utills.configureConnParamsFromZooKeeper(Utills.java:580)
    at org.apache.hive.jdbc.Utills.parseURL(Utills.java:391)
    at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.java:263)
    ... 5 more
    Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Tried all existing HiveServer2 uris from ZooKeeper.
    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.getServerHosts(ZooKeeperHiveClientHelper.java:132)
    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:172)
    ... 8 more

```

This issue occurs when you replicate at scale with an unbalanced cluster setup that has all the roles running on the Cloudera Manager host. As a result, Cloudera Manager ends up in a bottleneck situation because HS2 crashes, preventing further replication. The logs (/var/log/messages) indicate that there were a large number of Java processes running on the host, which exhausted the host's memory and triggered an Out Of Memory Killer process to stop HS2.

You can restart HS2 and reinitiate replication. However, the replication process may take a longer time to complete.

- Follow these recommendations when setting up your source and target clusters:
 - Ensure that no HS2, Hive metastore (HMS), DataNode, or NameNode roles are running on the Cloudera Manager host.
 - Ensure that you have multiple instances of HS2 and HMS roles on different nodes.
 - It is recommended that you have the HS2 and HMS roles on different nodes than the DataNodes or NameNodes.

These practices increase the possibility of Hive replication at scale completing successfully.

- Limit the scale of Hive replication by temporarily disabling replication policies.

CDPD-40730: Parquet change can cause incompatibility

Parquet files written by the parquet-mr library in CDP 7.1.9, where the schema contains a timestamp with no UTC conversion will not be compatible with older versions of Parquet readers. The effect is that the older versions will still consider these timestamps as they would require UTC conversions and will thus end up with a wrong result. You can encounter this problem only when you write Parquet-based tables using Hive, and tables have the non-default configuration `hive.parquet.write.int64.timestamp=true`.

None.

CDPD-41274: HWC + Oozie issue: Could not open client transport with JDBC Uri

Currently only Spark cluster mode is supported in the Oozie Spark Action with Hive Warehouse Connector (HWC).

Use Spark action in cluster mode.

```

<spark xmlns="uri:oozie:spark-action:1.0">
  ...
  <mode>cluster</mode>
  ...
</spark>

```

CDPD-26556: After an upgrade, querying a CTAS table under certain conditions might throw an exception

If you upgrade your Hive cluster from CDH 6 to CDP 7, create a CTAS table in the CDP cluster from a table you upgraded from CDH, you might see the following exception when you query the new table:

```
class org.apache.hadoop.io.IntWritable cannot be cast to class org.apache.hadoop.hive.serde2.objectinspector.StandardUnionObjectInspector$StandardUnion
```

This issue involves CDH-based tables having columns of complex types ARRAY, MAP, and STRUCT.

CDPD-45370: Cartesian join fails with an empty table

When the configuration `hive.tez.cartesian-product.enabled` is set to true, query fails if one of the tables involved in a Cartesian join has zero rows. This results in an `IllegalArgumentException` with the message `Number of items is 0. Should be positive`.

To address this issue, set `hive.tez.cartesian-product.enabled` to false before running the query.

Apache Jira: [\[HIVE-26582\]](#)

CDPD-23506: OutOfMemoryError in LLAP

Long running spark-shell applications can leave sessions in interactive Hiveserver2 until the Spark application finishes (user exists from spark-shell), causing memory pressure in case of a high number of queries in the same shell (1000+).

You must close spark-shell so that sessions are closed. Add the owner of the database or the tables as a user with read or read/write access to the tables directly.

CDPD-23041: DROP TABLE on a table having an index does not work

If you migrate a Hive table to CDP having an index, DROP TABLE does not drop the table. Hive no longer supports indexes ([HIVE-18448](#)). A foreign key constraint on the indexed table prevents dropping the table. Attempting to drop such a table results in the following error:

```
java.sql.BatchUpdateException: Cannot delete or update a parent row: a foreign key constraint fails ("hive"."IDXS", CONSTRAINT "IDXS_FK1" FOREIGN KEY ("ORIG_TBL_ID") REFERENCES "TBLS" ("TBL_ID"))
```

There are two workarounds:

- Drop the foreign key "IDXS_FK1" on the "IDXS" table within the metastore. You can also manually drop indexes, but do not cascade any drops because the IDXS table includes references to "TBLS".
- Launch an older version of Hive, such as Hive 2.3 that includes IDXS in the DDL, and then drop the indexes as described in [Language Manual Indexing](#).

Apache Issue: [Hive-24815](#)

CDPD-17766: Queries fail when using spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive.

Do not specify `spark.sql.hive.hiveserver2.jdbc.url.principal` in the JDBC URL to invoke Hive remotely.

Workaround: specify `principal=hive.server2.authentication.kerberos.principal` as shown in the following syntax:

```
jdbc:hive://<host>:<port>/<dbName>;principal=hive.server2.authentication.kerberos.principal;<otherSessionConfs>?<hiveConfs>#<hiveVars>
```

CDPD-13636: Hive job fails with OutOfMemory exception in the Azure DE cluster

Set the parameter `hive.optimize.sort.dynamic.partition.threshold=0`. Add this parameter in Cloudera Manager (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)

CDPD-10848: HiveServer Web UI displays incorrect data

If you enabled auto-TLS for TLS encryption, the HiveServer2 Web UI does not display the correct data in the following tables: Active Sessions, Open Queries, Last Max n Closed Queries

Technical Service Bulletins**TSB 2021-501: JOIN queries return wrong result for join keys with large size in Hive**

JOIN queries return wrong results when performing joins on large size keys (larger than 255 bytes). This happens when the fast hash table join algorithm is enabled, which is enabled by default.

Knowledge article

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2021-501: JOIN queries return wrong result for join keys with large size in Hive](#)

TSB 2023-702: Potential wrong result for queries with date partition filter for clusters in GMT+ timezone

In Cloudera Data Platform (CDP) Private Cloud Base 7.1.7 Service Pack (SP) 2 Cumulative Hotfix (CHF) 11, a fix was introduced in Hive Metastore (HMS) to address a parsing issue with date strings. This fix caused a regression in Hive clusters where the HMS time zone is set ahead of GMT for the following combination of tables and queries: a table that is partitioned on a DATE column and a SELECT query on that table containing a WHERE clause filter on the same DATE column. For such queries, during the partition pruning phase, the date string would be converted to a date without timezone and compared with the partition value retrieved by HMS. This causes wrong results (0 rows) because the date values do not match.

The regression was identified in CDP Private Cloud Base 7.1.7 SP2 CHF14, but it exists in CHF11 through CHF16 as well as on certain versions of 7.1.8 and 7.1.9.

This issue does not affect clusters where the time zones are behind GMT. For example, if the time zone of the cluster is set USA/Los Angeles, which is 8 hours behind GMT, a date '2023-10-02' will remain as '2023-10-02' after converting to GMT (adding 8 hours). On the other hand, using Asia/Hong Kong time as an example, which is 8 hours ahead of GMT, the same date would become '2023-10-01' after converting to GMT (subtracting 8 hours), which leads to the wrong results.

Upstream JIRA

[HIVE-27760](#)

Knowledge article

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2023-702: Potential wrong result for queries with date partition filter for clusters in GMT+ timezone](#)

Known Issues in Hue

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

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

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

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

1. Log in to Cloudera Manager as an Administrator.

2. Go to Clusters Hue Configurations and add the following key and value in the Hue Service Environment Advanced Configuration Snippet (Safety Valve) field:

Key: PYTHONPATH

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

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

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

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

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark queries in a batch mode launches a shell script Oozie job from Hue and this is not supported on a secure cluster.

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

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

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

None.

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

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

None.

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

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

None.

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

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

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

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

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

CDPD-43293: Unable to import Impala table using Importer

Creating Impala tables using the Hue Importer may fail.

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

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

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


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

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

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

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

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

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

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

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

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

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

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

Hue 4 web interface uses the unsafe-inline directive in its CSP header. As a result, the application server does not set the CSP header in its HTTP responses, and therefore does not benefit from the additional protection against potential cross-site scripting issues and other modern application vulnerabilities which a properly configured CSP may provide. This could lead to application vulnerability.

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

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

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

Manually restart the Hue service from Cloudera Manager.

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

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

None

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

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

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

1. Log in to Cloudera Manager as an Administrator.
2. Obtain the host name of the Knox Gateway by going to Clusters Knox service Instances .

3. Go to **Clusters Hue service Configuration** and search for the **Knox Proxy Hosts** field.



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

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

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

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

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

```
[search]
# URL of the Solr Server
solr_url=http://[***HOSTNAME***]:[***PORT***]/solr/
```

For example:

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

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

CLR-72251: Invalid S3 URI error while accessing S3 bucket

The **Hue** Load Balancer merges the double slashes (//) in the **S3** URI into a single slash (/) so that the URI prefix `/filebrowser/view=S3A://` is changed to `/filebrowser/view=S3A:/`. This results in an error when you try to access the **S3** buckets from the **Hue** File Browser through the port 8889.

The **Hue** web UI displays the following error: “Unknown error occurred”.

The **Hue** server logs record the “ValueError: Invalid S3 URI: S3A” error.

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.

You should be able to load the **S3** browser from both 8888 and 8889 ports.

Alternatively, you can use the **Hue** server port 8888 instead of the load balancer port 8889 to resolve this issue.

CLR-72255: Error while rerunning Oozie workflow

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

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

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

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

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

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

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

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

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

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

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

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

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

Query history displays "Result Expired" message

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

None.

Unsupported features

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

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

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

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

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

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

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

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

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

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



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

Technical Service Bulletins

TSB 2023-704: File corruption when downloading files larger than 1 MB from ABFS with Hue File Browser

An issue within the upstream Apache Hue (Hue) application results in file corruption when downloading files larger than 1MB files from Azure Blob Filesystem (ABFS) using the Hue File Browser. Only the downloaded files are affected by this issue, the source files remain intact.

Knowledge article

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2023-704: File corruption when downloading files larger than 1 MB from ABFS with Hue File Browser](#)

TSB 2023-703: Risk of Data Loss when using Hue S3 File Browser

There is a risk of losing data when moving one or more files or folders in Amazon S3 storage with the Hue File Browser. When the user selects the destination folder in the modal window the following scenarios can occur:

1. If the user selects the same destination folder as the source folder, and clicks the Move button, the selected files will be permanently deleted.
2. If the user selects a different destination folder from the source folder and clicks the Move button before the User Interface (UI) has completely loaded (the loading is indicated by a spinner), the action could lead to the following outcomes:
 - a. If the previously visible destination folder was the same as the source folder, the file will be permanently deleted.
 - b. If the previously visible destination folder was different from the source folder, the file(s) will be moved to the previously visible destination folder and not to the intended destination folder.

Knowledge article

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2023-703: Risk of Data Loss when using Hue S3 File Browser](#)

Known Issues in Iceberg

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

CDPD-75422: Impala schema case sensitivity issue with workaround

Impala's schema is case insensitive, causing errors with mixed case schema elements created through Spark during predicate pushdown.

- Create tables through Impala to ensure lower case schema.
- Avoid upper case in Spark: Do not use upper case when creating tables through Spark.
- Fix existing tables: Use ALTER TABLE to rename upper case columns:

```
ALTER TABLE `database`.`iceberg_table` CHANGE COLUMN ID id string;
```

Apache Jira: [IMPALA-13463](#)

Known Issues in Impala

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

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

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

The following services support DDL statements:

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

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

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

CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the `--tls_ciphersuites` startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.

Ensure that the list in the `--tls_ciphersuites` startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

Apache Jira: [IMPALA-14625](#)

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work

If you try to run `impala-shell` on RHEL 9 by setting the default python executable available in `PATH` to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, `IMPALA_PYTHON_EXECUTABLE=python3`.

CDPD-42958: After upgrading the CDH 7.1.9 from CDH 6.x, under certain conditions you cannot insert data into a table

Under the following conditions, after upgrading from CDH 6.x to CDH 7.1.9 you cannot insert data into a table from Impala:

- On CDH 6.x, you created a database with Impala in a user specified HDFS location.
- Using Hive, you then created a table in the database.

Under these conditions, the database and table are stored in the user-specified HDFS directory. After upgrading, the HDFS directory of the table is read-only for Impala. Consequently, from Impala you cannot insert new data into the table because Impala does not have write permission on the HDFS directory.

Workaround: To resolve this issue, use either one of the following workarounds:

- Using the Ranger Web UI, in the policy repository `cm_hdfs`, grant the user 'impala' write permission on the directory where the table resides.
- Enter the following command to grant write permission to user 'impala' on the HDFS directory where the table resides.

```
hdfs dfs -setfacl -m default:user:impala:rwx <HDFS directory>
```

CDPD-80019: Impala fails to start due to large catalog-update topic

Impala coordinators fail to start if the catalog-update statestore topic exceeds 2GB, hitting the Thrift 0.16.0 message size limit (`thrift_rpc_max_message_size`). This prevents the coordinator from processing the update.

- Enable local catalog mode: `--use_local_catalog=true`
- In Cloudera Manager, go to Impala Configuration search for "Catalog Server Command Line Argument Advanced Configuration Snippet (Safety Valve)" and click 
- Add property name and value:

```
Name : --catalog_topic_mode
Value: minimal
```

- Save changes.

Apache Jira: [IMPALA-12350](#)

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 HDF
S file hdfs://nameservice1/warehouse/tablespace/managed/hive/<da
tabase>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException:
File does not exist: /warehouse/tablespace/managed/hive/<data
base>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

You must use a single Knox node to access Impala UI.

Impala api calls via 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-21828: Multiple permission assignment through grant is not working

None

Problem configuring masking on tables using Ranger

The following Knowledge Base article describes the behavior when we configure masking on tables using Ranger. This configuration works for Hive, but breaks queries in some scenarios for Impala.

For a workaround, see the following Knowledge Base article: [ERROR: "AnalysisException: No matching function with signature: mask\(FLOAT\)" when Impala jobs fail with the following error with signature: mask\(FLOAT\)](#)

IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXXXX"

IMPALA-532: Impala should tolerate bad locale settings

If the LC_* environment variables specify an unsupported locale, Impala does not start.

Add LC_ALL="C" to the environment settings for both the Impala daemon and the Statestore daemon.

Avro Scanner fails to parse some schemas

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

IMPALA-691: Process mem limit does not account for the JVM's memory usage

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the impalad daemon.

To monitor overall memory usage, use the `top` command, or add the memory figures in the Impala web UI `/memz` tab to JVM memory usage shown on the `/metrics` tab.

IMPALA-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Remove trailing semicolon from the Avro schema.

IMPALA-1652: Incorrect results with basic predicate on CHAR typed column

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the `RPAD()` function to blank-pad literals compared with `CHAR` columns to the expected length.

IMPALA-1821: Casting scenarios with invalid/inconsistent results

Using a `CAST()` function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

IMPALA-2005: A failed CTAS does not drop the table if the insert fails

If a `CREATE TABLE AS SELECT` operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

Drop the new table manually after a failed CREATE TABLE AS SELECT

IMPALA-2422: % escaping does not work correctly when occurs at the end in a LIKE clause

If the final character in the RHS argument of a LIKE operator is an escaped \% character, it does not match a % final character of the LHS argument.

None

IMPALA-2603: Crash: impala::Coordinator::ValidateCollectionSlots

A query could encounter a serious error if includes multiple nested levels of INNER JOIN clauses involving subqueries.

None

IMPALA-3094: Incorrect result due to constant evaluation in query with outer join

An OUTER JOIN query could omit some expected result rows due to a constant such as FALSE in another join clause. For example:

```
explain SELECT 1 FROM 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 |
+-----+
```


IMPALA-3509: Breakpad minidumps can be very large when the thread count is high

The size of the breakpad minidump files grows linearly with the number of threads. By default, each thread adds 8 KB to the minidump size. Minidump files could consume significant disk space when the daemons have a high number of threads.

Add `-\-minidump_size_limit_hint_kb=size` to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then 2 KB per thread after that.) The minidump file can still grow larger than the "hinted" size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

IMPALA-4978: Impala requires FQDN from hostname command on Kerberized clusters

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the `gethostname()` system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a Kerberized cluster.

Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname` only returns the short name, pass the command-line flag `##hostname=FULLY_QUALIFIED_DOMAIN_NAME` in the startup options of all Impala-related daemons.

IMPALA-6671: Metadata operations block read-only operations on unrelated tables

Metadata operations that change the state of a table, like `COMPUTE STATS` or `ALTER RECOVER PARTITIONS`, may delay metadata propagation of unrelated unloaded tables triggered by statements like `DESCRIBE` or `SELECT` queries.

None

IMPALA-7072: Impala does not support Heimdal Kerberos

None

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.

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.

Some of the unresolved issues include:

- IMPALA-6841
- IMPALA-635

Known Issues in Kafka

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

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

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

The following services support DDL statements:

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

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

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

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-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker) in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Save Changes > Restart SMM.

The offsets.topic.replication.factor property must be less than or equal to the number of live brokers

The offsets.topic.replication.factor broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a GROUP_COORDINATOR_NOT_AVAILABLE error until the cluster size meets this replication factor requirement.

None

Requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true

The first few produce requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true.

Increase the number of retries in the producer configuration setting retries.

Performance degradation when SSL Is enabled

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with ssl.secure.random.implementation = SHA1PRNG. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

OPSAPS-43236: Kafka garbage collection logs are written to the process directory

By default Kafka garbage collection logs are written to the agent process directory. Changing the default path for these log files is currently unsupported.

None

RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transactional method because we are in an error state
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeFailWithError(TransactionManager.java:1125)
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeAddPartition(TransactionManager.java:442)
    at org.apache.kafka.clients.producer.KafkaProducer.doSend(KafkaProducer.java:1000)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:914)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:800)
    .
    .
    .
    Caused by: org.apache.kafka.common.errors.ClusterAuthorizationException: Cluster authorization failed.
```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting `enable.idempotence` to `false`.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

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

The Debezium Db2 Source connector does not support the evolution (updates) of schemas. In addition, schema change events are not emitted to the schema change topic if there is a change in the schema of a table that is in capture mode. For more information, see [DBZ-4990](#).

None.

CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy compression. The HDFS Stateless Sink connector is only affected if the Compression Codec or Compression Codec for Parquet properties are set to SNAPPY.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: org.apache.hadoop.util.NativeCodeLoader.buildSupportsSnappy()
```

```
Failed to write to HDFS due to java.lang.RuntimeException: native snappy library not available: this version of libhadoop was built without snappy support.
```

Download and deploy missing libraries.



Important: Ensure that you complete steps 1-11 on all Kafka Connect hosts. Additionally, ensure that the advanced configuration snippet in step 12 is configured for all Kafka Connect role instances.

1. Create the /opt/nativelibs directory.

```
mkdir /opt/nativelibs
```

2. Change the owner to kafka.

```
chown kafka:kafka /opt/nativelibs
```

3. Locate the directory containing the Hadoop native libraries and copy its contents to the directory you created.

```
cp /opt/cloudera/parcels/CDH/lib/hadoop/lib/native/* /opt/nativelibs
```

4. Verify that libsnpappy.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-69481: Some Kafka Connect metrics missing from Cloudera Manager due to conflicting definitions

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

Contact Cloudera support for a workaround.

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

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Unsupported Features

The following Kafka features are not supported in Cloudera Data Platform:

- Only Java and .Net based clients are supported. Clients developed with C, C++, Python, and other languages are currently not supported.
- The Kafka default authorizer is not supported. This includes setting ACLs and all related APIs, broker functionality, and command-line tools.
- SASL/SCRAM is only supported for delegation token based authentication. It is not supported as a standalone authentication mechanism.
- Kafka KRaft in this release of Cloudera Runtime is in technical preview and does not support the following:
 - Deployments with multiple log directories. This includes deployments that use JBOD for storage.
 - Delegation token based authentication.
 - Migrating an already running Kafka service from ZooKeeper to KRaft.
 - Atlas Integration.

Limitations

Collection of Partition Level Metrics May Cause Cloudera Manager's Performance to Degrade

If the Kafka service operates with a large number of partitions, collection of partition level metrics may cause Cloudera Manager's performance to degrade.

If you are observing performance degradation and your cluster is operating with a high number of partitions, you can choose to disable the collection of partition level metrics.



Important: If you are using SMM to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both SMM and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, SMM will not be able to display information about partitions. In addition, Cruise Control will not operate properly.

Complete the following steps to turn off the collection of partition level metrics:

1. Obtain the Kafka service name:
 - a. In Cloudera Manager, Select the Kafka service.
 - b. Select any available chart, and select Open in Chart Builder from the configuration icon drop-down.
 - c. Find \$SERVICENAME= near the top of the display.

The Kafka service name is the value of \$SERVICENAME.
2. Turn off the collection of partition level metrics:
 - a. Go to Hosts Hosts Configuration .
 - b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

Enter the following to turn off the collection of partition level metrics:

```
[KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_entity_update_enabled=false
```

Replace [KAFKA_SERVICE_NAME] with the service name of Kafka obtained in step 1. The service name should always be in lower case.

- c. Click Save Changes.

Known Issues in Kerberos

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

Known issues in Key Trustee Server

Learn about the known issues in Key Trustee Server, the impact or changes to the functionality, and the workaround.

KTS setup on FIPS cluster gets stuck if KTS is in HA

Key Trustee Server setup gets stuck if KTS is in HA for Centos 7.9 in FIPS mode.

Create Date is not fetched from Key Trustee Server

During key migration from KTS to Ranger KMS , create date is not fetched from KTS.

Known Issues in Knox

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

CDPD-79963: Knox service might fail due to JARs picked up from the /usr/share/java folder

Knox service might fail due to Java Archive (JAR) files picked up from the /usr/share/java folder.

None.

CDPD-67478: Custom topologies cannot be deleted

You cannot delete custom topologies that were created.

- Edit the custom topology in Cloudera Manager, removing all service declarations and leaving only the providerConfigRef configuration.
- Refresh the Knox configuration.

Knox Issue with JDK version

jdk-1.8.0_391 is not supported.

Cloudera recommends using Cloudera supported JDKs.

CDPD-61088: When downgrade is performed from CDP 7.1.9 to CDP 7.1.7 SP2, Knox may fail to start

When you downgrade from CDP 7.1.9 to CDP 7.1.7 SP2, Knox might fail to start with the following error message:

Failed to start gateway: org.apache.knox.gateway.services.ServiceLifecycleException: Keystore was not loaded properly - the provided password may not match the password for the keystore. org.apache.knox.gateway.services.ServiceLifecycleException: Keystore was not loaded properly - the provided password may not match the password for the keystore.

Remove the faulty credential store and restart Knox.

CDPD-60996: When downgrade is performed from CDP 7.1.9 to CDP 7.1.7 SP2, Knox is unable to connect to Cloudera Manager.

Restart Knox service after the downgrade process completes.

CDPD-28431: Intermittent errors can be potentially encountered when Impala UI is accessed from multiple Knox nodes.

You must use a single Knox node to access Impala UI.

CDPD-3125: Logging out of Atlas does not manage the external authentication

At this time, Atlas does not communicate a log-out event with the external authentication management, Apache Knox. When you log out of Atlas, you can still open the instance of Atlas from the same web browser without re-authentication.

To prevent additional access to Atlas, close all browser windows and exit the browser.

CDPD-22785: Improvements and issues needs to be addressed in convert-topology Knox CLI command

None

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

Manually add cdp-proxy-token to the Knox policy, using Ranger Admin Web UI.

1. Log in to Cloudera Manager Ranger Admin Web UI, as a Ranger administrator.
2. On Ranger Admin Web UI Service Manager Resource Knox, click cm_knox.
3. In Knox Policies, open the CDP Proxy UI, API and Token policy.
4. In Knox Topology*, add cdp-proxy-token.
5. Click Save.
6. Restart Ranger.

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

KNOX does not send neither the Authentication header nor the hadoop.auth cookie. Because of this, the SMM UI returns an HTTP 401 response and sets the www-authenticate: Negotiate header. After this, KNOX still does not send the Authentication header. This results in the SMM UI being inaccessible through Knox.

You can access the SMM UI directly in Cloudera Manager at Clusters SMM Streams Messaging Manager Web UI and log in using the Kerberos username and password.

On FIPS cluster, Knox-Impala connection is failing with SSL error code 5

The Knox-Impala connection is failing with SSL error code 5 on FIPS clusters. This prevents accessing Impala through Knox.

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

The health test for Knox Gateway fails if the gateway.client.auth.needed parameter is set to true. Environments using the "curl" call are impacted. The curl call from CM is not specifying any certificate (store) while Knox is configured to require one.

To avoid the startup failure, use gateway.client.auth.wanted = true.

CDPD-71751: Creation of alias from the Cloudera Manager UI fails on FIPS

Users attempting to create aliases through the Cloudera Manager UI face issues in FIPS.

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

1. ssh to Knox host.
2. Export these directories: export KNOX_GATEWAY_DATA_DIR="/var/lib/knox/gateway/data"; export KNOX_GATEWAY_CONF_DIR="/var/lib/knox/gateway/conf"
3. Set the FIPS-specific options for the Knox CLI:

```
export KNOX_CLI_MEM_OPTS="- --add-exports=java.base/sun.security.provider=bctls --add-exports=java.base/sun.security.provider=com.safelogic.cryptocomply.fips.core --add-modules=com.safelogic.cryptocomply.fips.core --add-modules=bctls --module-path=<BCTLS_JARS_DIR> -Dcom.safelogic.cryptocomply.fips.approved_only=true"
```

<BCTLS_JARS_DIR> is the directory containing the SafeLogic bctls and fips core jar files.

4. Run the following command to create the alias: /opt/cloudera/parcels/CDH/lib/knox/bin/knoxcli.sh create-alias <ALIAS_NAME> <ALIAS_VALUE>
5. Verify the addition using /opt/cloudera/parcels/CDH/lib/knox/bin/knoxcli.sh list-alias.

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

Known Issues in Kudu

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

- Kudu HMS Sync is disabled and is not yet supported

You get "The user 'kudu' is not part of group 'hive' on the following hosts: " warning by the Host Inspector

If you are using fine grained authorization for Kudu, and you are also using Kudu-HMS integration with HDFS-Sentry sync, then you may get the "The user 'kudu' is not part of group 'hive' on the following hosts: " warning while upgrading.

Workaround: Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

Known Issues in Livy

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

Known Issues

No known issues identified in Livy.

Known Issues in Navigator Encrypt

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

Failed to start navencrypt-mount.service

After installing NavEncrypt on SLES, if the command "systemctl status navencrypt-mount" fails with the error: "Failed to start navencrypt-mount.service: Unit navencrypt-mount.service failed to load: No such file or directory", the systemd files need to be reloaded..

Run the command:

```
(sudo) systemctl daemon-reload
```

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

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 Oozie

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

OPSAPS-74399: Missing jars from Spark 3 ShareLib when using with CDS 3.5

The `spark-common-utils_2.12-3.5.4.3.5.7191000.0-XX.jar` and `spark-sql-api_2.12-3.5.4.3.5.7191000.0-XX.jar` files are missing from the Spark 3.5.4 sharelib directory causing the Spark 3 actions to fail.

Perform the troubleshooting steps as described in [Troubleshoot Spark 3 ShareLib compatibility behaviour](#).

CDPD-41274: Hive Warehouse Connector (HWC) and Oozie issue: Can not open client transport with JDBC URI

Currently, only Spark cluster mode is supported in the Oozie Spark action with Hive Warehouse Connector (HWC).

```
Use Spark action in cluster mode.
    <spark xmlns="uri:oozie:spark-action:1
.0">
        ...
    <mode>cluster</mode>
    ...
</spark>
```

CDPD-26975: Using the ABFS / S3A connectors in an Oozie workflow where the operations are "secured" may trigger an IllegalArgumentException with the error message java.net.URISyntaxException: Relative path in absolute URI.

Set the following XML configuration in the Datahub cluster's Cloudera Manager:

1. In the Cloudera Manager Admin Console, go to the Oozie service.
2. Go to the Configuration tab.
3. In the Oozie Server Advanced Configuration Snippet (Safety Valve) for `oozie-site.xml` field, set the following:

Set the following if you are using Amazon S3:

```
<property>
pAccessorService.fs.s3a</name>
tmp/s3a</value>
    <name>oozie.service.Hadoop
    <value>fs.s3a.buffer.dir=
```

```
</property>
```

Set the following if you are using ABFS:

```
<property>
  <name>oozie.service.HadoopAccessorService.fs.abfs</name>
  <value>fs.azure.buffer.dir=/tmp/abfs</value>
</property>

<property><name>oozie.service.HadoopAccessorService.fs.abfss</name>
  <value>fs.azure.buffer.dir=/tmp/abfss</value>
</property>
```

4. Enter a Reason for change, and then click Save Change to commit the changes.
5. Restart the Oozie service.

OOZIE-3549: Oozie fails to start when Cloudera Manager 7.x is used with Cloudera Runtime 6.x and Java 11 because Oozie does not set the trust-store password.

The issue is fixed in OOZIE-3549 and is already included in CDP 7.x but not in CDH 6.x. If you are on CDH 6.x and want to upgrade to Java 11 or your Cloudera Manager version to 7.x then you must request for a patch.

Oozie jobs fail (gracefully) on secure YARN clusters when the JobHistory server is down.

If the JobHistory server is down on a YARN (MRv2) cluster, Oozie attempts to resubmit a job three times by default. If the job fails, Oozie automatically puts the workflow in a SUSPEND state.

When the JobHistory server is running again, use the `resume` command to inform Oozie to continue the workflow from the point at which it stopped.

Unsupported Feature

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

- Non-support for Apache Pig action (CDPD-1070)
- Conditional coordinator input logic process

Cloudera does not support using Derby database with Oozie. You can use it for testing or debugging purposes, but Cloudera does not recommend using it in production environments. This may cause failures while upgrading from CDH to CDP.

Known Issues in Ozone

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

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: 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.

ENGESC-26990: Standalone Ozone deployment in a cluster is not supported.

A cluster with Ozone but without HDFS is not supported. Setting the `fs.defaultFs` parameter to Ozone is not certified yet.

None.

The Prometheus binary is not available in CDP Private Cloud Base 7.1.9 SP1 for the Ubuntu operating system.

You can install Prometheus separately and specify the path to the parent directory, for example `/usr/bin`, in the `prometheus.location` parameter in Ozone.

In CDP Private Cloud Base 7.1.9 SP1, there is a missing block location in the output of Ozone debug chunkinfo command for EC bucket.

None.

In CDP Private Cloud Base 7.1.9 SP1 and CDP Private Cloud Base 7.1.7 SP3, there is an intermittent key put failure after stopping the OM leader.

None.

In CDP Private Cloud 7.1.9 SP1, Recon certificates are not available after running the `ozone admin cert list` command. This issue is present when upgrading from a lower version of CDP to CDP Private Cloud 7.1.9 SP1.

1. Ensure that the Recon user is included in the `hdds.security.client.datanode.container.protocol.acl` configuration in the `hadoop-policy.xml` file under `/var/run/cloudera-scm-agent/process/*-ozone-STORAGE_CONTAINER_MANAGER/ozone-conf`.
2. If the Recon user is not listed, SCM rejects Recon's request for certificates. This step must be completed before proceeding to the next step.
3. Remove the keys folder in the Recon metadata directory and restart the cluster. This action results in new keys and certificates being generated from SCM.

On upgrading the cluster from CDP Private Cloud Base 7.1.7 SP3 to CDP Private Cloud Base 7.1.9 SP1, you cannot create snapshots on a pre-upgrade volume or bucket and `PERMISSION_DENIED` error is displayed. This is because from CDP Private Cloud Base 7.1.7 SP3, complete Kerberos owner name with domain is used. However, in CDP Private Cloud Base 7.1.9 SP1 the Kerberos owner short name is used.

You must update the Kerberos owner name for the bucket with the user short name and then create snapshot by running the following commands:

- `ozone sh bucket update --user={user_short_name} {volume}/{bucket}`
- `ozone sh snapshot create {volume}/{bucket} {snapshot_name}`

CDPD-70409: Recon Overview Page UI fails to load if Recon Solr health displays an error. This is because the Solr service is not enabled on the cluster.

You must manually enable the Solr service on the cluster:

1. Log in to Cloudera Manager UI.
2. Navigate to Clusters.
3. Select the Ozone service.
4. Go to Configurations.
5. Search for `solr_service` and select the option.
6. Click Save.

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 CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and if Ozone is in the Non-HA environment, an exception message is observed during the finalization of the Ozone upgrade.

During the finalization of the upgrade, a `ClassNotFoundException` message for `org.cloudera.log4j.redactor.RedactorAppender` class may be displayed. The error message can be ignored because the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

CDPD-68951: In CDP Private Cloud Base 7.1.9 CHF2 version and lower, the command `ozone sh key list <bucket_path>` displays the `isFile` flag in a key's metadata as false even when the key is a file. This issue is rectified in CDP Private Cloud Base 7.1.9 CHF3. However, the pre-existing (pre-upgrade) key's metadata cannot be changed.

None

CDPD-67771: 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.

CDPD-64398: In SCM with non-HA configuration, Secret key manager does not initialise. Hence, the startup of OM and DataNode fails because it cannot get a secret key. This key is used when security is enabled and used in Block token and container token verification while communicating between the Ozone client, OM, and Datanode.

This issue does not occur in SCM with HA configuration.

You must force exit from the safe mode for SCM. This triggers the initialization of the Secret key manager. Below are the options:

- exit safe mode manually whenever SCM is started
- disable safe mode by setting `hdds.scm.safemode.enabled=false` in safety valve for SCM configuration

Impact of disabling safe mode and the purpose of safe mode:

- Write should not fail once the cluster is out of safe mode
- Read of existing data should not fail after the cluster is out of safe mode
- Unnecessary re-replication should be avoided during cluster restart

So disabling safe mode does not have any major impact on Ozone.

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.

CDPD-52412: keyManagerImpl#listStatus exceeds the maximum RPC length.

To reduce the payload, you must increase the part size.

OPSAPS-68159: If you did not deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Error when distributing parcel to host, Ignoring non-compliant parcel manifest error is displayed after Cloudera Manager is upgraded to 7.11.3.

If you encounter the error, perform the following steps:

1. Deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.11.3.

2. Restart the cluster with a delay of 10 minutes.
3. Continue to upgrade CDH 7.1.8 to CDH 7.1.9.

OPSAPS-68159: If you did not deactivate the Ozone parcel 718.2.x on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Ozone roles during the CDH 7.1.8 upgrade to CDH 7.1.9.

If you encounter the error, perform the following steps:

1. Deactivate the Ozone parcel 718.2.x.
2. Restart the Ozone service.
3. Perform Finalize Upgrade for the Ozone service.

Step result: The Ozone roles will be displayed in green.

CDPD-60989: The packaging version for Apache Ozone points to the 1.2.0 older version. This is a version string problem and not a packaging issue. The version of the Apache Ozone binary is closest to 1.4.0.

None. This only affects the JAR versioning.

CDPD-60366: Native library loader fails when system property native.lib.tmp.dir is not set. It fails because the library is copied to / instead of the cwd.

Setting native.lib.tmp.dir to a certain path like /tmp should solve the issue. To set this system property, you must add -Dnative.lib.tmp.dir=/tmp to ozone_java_opts in Cloudera Manager configuration of the Ozone cluster.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-60598: DataNode can consume up to the JVM heap size worth of direct memory buffers under specific failure scenarios. This can lead to a DataNode crash. This is due to Netty allocating and not freeing the direct memory buffers.

Restarting DataNode will resolve this issue.

CDPD-60466: The Ozone client is missing in the Cloudera Manager's Service Monitor node. This causes Ozone Canary's health check to fail.

Ensure you install Ozone CLI on the same node as Cloudera Manager's Service Monitor.

CDPD-60012: The Ozone SCM may be stuck in safe mode after upgrade to 7.1.9. This is due to a bug in how the SCM accounts for open storage containers.

The workaround is to temporarily set the Ozone configuration hdds.scm.safemode.threshold.pct to a lower value like 0.90 and restart the SCM.

CDPD-59679: In some scenarios, excess UNHEALTHY replicas of EC containers may not be removed from the cluster.

None

CDPD-50116: Topology-aware reads can provide better performance by routing applications to the closest replica of a file, if available, on the same physical node or same rack. However this feature is disabled by default in CDP 7.1.9.

This feature can be enabled by setting ozone.network.topology.aware.read to true using Cloudera Manager safety valve.

CDPD-57165: The DataNode disk usage thread may abort due to Ratis log tailing.

Use org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory to calculate disk usage. Configure in the ozone-site.xml property: hdds.datanode.du.factory value: org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory

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.

CDPD-55513: Hive external table replication policy does not migrate Hive tables in HDFS storage to Ozone.

To replicate data from HDFS to Ozone, see [Migrating your data from HDFS to Ozone](#). You can use HMS Mirror to replicate the metadata.

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-64097: Ozone service restart failed at SCM

Stopping SCM service using Cloudera Manager can sometimes timeout and need a retry. The Cloudera Manager API waits for 90 seconds which is not sufficient under certain circumstances.

Retry the shutdown using Cloudera Manager if the SCM still shows as running after a refresh of the service status.

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-filesystem-hadoop3.jar, it is likely due to the presence of the user hdfs on the node prior to CDP parcel activation. As a result, the Cloudera Manager agent skips the initialization related to HDFS, leading to the non-creation of the /var/lib/hadoop-hdfs directory.

Create the directory “/var/lib/hadoop-hdfs” on all nodes followed by the deactivation and activation of the CDP parcel (deactivate and activate the Ozone parcel instead, in case Ozone parcel is used).

CDPD-50447: When SCM High Availability is enabled, each of the SCM web UIs report the host of the web UI as the leader of HA, and the other two as followers. This gives wrong information

Correct output is available by running the `ozone admin scm roles --service-id=<ID>` command.

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](#).

OPSAPS-62327: In an Ozone cluster without any gateway roles, Ozone is unable to deploy client configurations and displays the ConfigGenException error.

You must add the Ozone gateway roles to the cluster.

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-36389: The configurations `datanodes.involved.max.percentage.per.iteration` and `size.moved.max.per.iteration` are meant to limit the max number of DataNodes that is involved and maximum size that can move in an iteration. This bug causes balancer to stop an iteration when it is two DataNodes or 1 Container size (5GB) away from hitting these limits. However, these DataNodes can again be considered for balancing in the next iteration. This means the cluster is balanced after enough iterations, albeit a bit slowly. This bug is apparent in small clusters of around four DataNodes where the DataNodes can be either the source or target for a lot of moves but the iteration gets stopped when three DataNodes have been involved. It takes a higher number of iterations to eventually balance this cluster. While this is a performance issue, it does not prevent balancer from ultimately balancing the cluster.

Increase the speed for balancing by decreasing the interval between each iteration using the configuration `balancing.iteration.interval`. The value of this configuration must be greater than `hdds.datanode.du.refresh.period`. `size.moved.max.per.iteration` can be increased to allow more data to move in one iteration.

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-35141: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive.ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2> (state=08S01,code=40000) java.sql.SQLException: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive.ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2>. This issue does not occur if the source and target buckets are different in Hive queries. Currently, copying across the same bucket is only supported.

Avoid different buckets in source and target path. Ensure that hive.exec.stagingdir is unset or set to the same bucket as the warehouse so that staging directories are created in the same bucket as the warehouse.

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-41184: With Legacy buckets, FileSystem Optimised is not interoperating with the Ozone shell command. The cause is that the directory key entry in the DB KeyTable is stored as "dir1/" with trailing slash. But while performing the described operation, Ozone shell (o3://) is normalizing the given path and removing the trailing slash "/" from it. That results in KEY_NOT_FOUND exception.

There are three workarounds:

- Use FileSystem API to Delete the Directories rather than Shell-Command API.
- Use FSO buckets instead of Legacy Buckets. As in FSO, you can create Intermediate Directories and Delete Directories using the Ozone shell commands.
- Disable and set the configuration ozone.om.enable.filesystem.pathsflag to false in order to delete the directories. This is generally not a preferred workaround because the cluster must be restarted to pick up the new changes.

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-12542: Mount of Ozone filesystem with the help of FUSE fails.

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-42897: EC writes are failing with "No enough datanodes to choose" after EC replication config is set globally.

EC writes start failing when large number of pipelines are created as a result of multiple EC configs with different chunk sizes used to write keys.

If standard EC configs (that is, rs-3-2-1024k) are used to write keys, number of pipelines created per DataNode is limited to five and this issue is not seen with standard EC configs.

The recommendation is not to create too many random chunk sizes. It is configurable because, users can decide based on their workload. But do not to have separate chunk sizes for each file.

CDPD-41539: The "No such file or directory" imessage is displayed when EC file is read using older of client.

You must upgrade the client before trying to read the key: `vol1/ecbuck1/1GB_ec`.

CDPD-40560: Filesystem Operations using Hadoop S3a connector on a FILE_SYSTEM_OPTIMIZED bucket fails. `org.apache.hadoop.ozone.om.exceptions.OMException: Unable to get file status: volume: s3v bucket: fso key: test/`

Do not run Hadoop S3a commands on a FILE_SYSTEM_OPTIMIZED bucket. Use OBJECT_STORE bucket layouts.

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.

CDPD-43432: Ozone Service in fault state in DataNode - Long Running setup.

Upgrade RocksDB to the latest version.

OPSAPS-63999: In the newly installed cluster, the Finish upgrade option is clickable.

None

OPSAPS-64648: Failed to start Ozone node using Cloudera Manager if default log path `/var/log/hadoop-ozone` does not exist. If this path does not exist, any Ozone nodes (for example SCM or DataNode) restart will fail.

Run the following command `sudo -u hdfs mkdir -p /var/log/hadoop-ozone` or replace hdfs with the user Ozone roles that are running.

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.

Technical Service Bulletins**TSB 2025-820: Potential Data Integrity Issues Found in Ozone**

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

Knowledge article

For the latest update on this issue see the corresponding Knowledge 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

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.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

Known Issues in Parquet

There are no known issues for Parquet in Cloudera Runtime 7.1.9 SP1.

Known Issues in Phoenix

There are no known issues for Phoenix in Cloudera Runtime 7.1.9 SP1.

Known Issues in Queue Manager

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

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.

Known Issues in Ranger

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

OPSAPS-75673: Wrong enablement of Ranger RMS Database Full Sync command

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

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.

CDPD-66092: Upgrade from CDP-7.1.8 to 7.1.9 fails, Ranger shows no policies at all, due to Java patches not being applied during upgrade. This causes rolling restart of services to fail.

Skip applying any of the following java patches NOT applicable to the underlying environment. In other words, do not apply patches for a service definition that does not appear in ranger database.

Table 1: Java Patches and Related Service Definitions

Java patch	Service-definition
PatchForHiveServiceDefUpdate_J10027.java	Hive
PatchForTagServiceDefUpdate_J10028.java	Tag
PatchForHBaseServiceDefUpdate_J10035.java	HBase
PatchForHdfsAddChainedPluginProvider_J10038.java	HDFS
PatchForHdfsRemoveChainedPluginProvider_J10039.java	HDFS
PatchForOzoneServiceDefUpdate_J10041.java	Ozone
PatchForOzoneServiceDefConfigUpdate_J10051.java	Ozone
PatchForOzoneServiceDefUpdate_J10057.java	Ozone

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-56803: When there is no existing policy for user and a revoke request comes from hbase, then will get this error.

```
hbase:001:0> revoke 'hrt_11'
```

```

ERROR: org.apache.hadoop.hbase.coprocessor.CoprocessorException:
HTTP 400 Error: processSecureRevokeRequest processing failed
    at org.apache.ranger.authorization.hbase.RangerAuthorizationCopr
cessor.preRevoke(RangerAuthorizationCoprocessor.java:1309)
    at org.apache.ranger.authorization.hbase.RangerAuthorizationCopr
cessor.preRevoke(RangerAuthorizationCoprocessor.java:1128)
    at org.apache.hadoop.hbase.master.MasterCoprocessorHost$162.ca
ll(MasterCoprocessorHost.java:1857)
    at org.apache.hadoop.hbase.master.MasterCoprocessorHost$162.ca
ll(MasterCoprocessorHost.java:1854)
    at org.apache.hadoop.hbase.coprocessor.CoprocessorHost$Observer0
perationWithoutResult.callObserver(CoprocessorHost.java:558)
    at org.apache.hadoop.hbase.coprocessor.CoprocessorHost.execOpe
ration(CoprocessorHost.java:631)
    at org.apache.hadoop.hbase.master.MasterCoprocessorHost.preRev
oke(MasterCoprocessorHost.java:1854)
    at org.apache.hadoop.hbase.master.MasterRpcServices.revoke(Mas
terRpcServices.java:2740)
    at org.apache.hadoop.hbase.shaded.protobuf.generated.MasterPr
otos$MasterService$2.callBlockingMethod(MasterProtos.java)
    at org.apache.hadoop.hbase.ipc.RpcServer.call(RpcServer.java:387
)
    at org.apache.hadoop.hbase.ipc.CallRunner.run(CallRunner.java
:139)
    at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExe
cutor.java:369)
    at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExe
cutor.java:349)

```

None

CDPD-56741: Improvement in log message when jwtauth not used

The following exception is printed at startup only and it is not cluttering logs:

```

2023-05-30 06:18:40,127 ERROR org.apache.ranger.rms.security.RMS
JwtAuthFilter:
quasar-pibgzl-1.quasar-pibgzl.root.hwx.site-startStop-1]:
Failed to initialize Ranger RMS JWT Auth Filter.
java.lang.Exception: RangerJwtAuthHandler:
Mandatory configs ('jwks.provider-url' & 'jwt.public-key') are m
issing, must provide atleast one.
    at org.apache.ranger.authz.handler.jwt.RangerJwtAuthHandler.in
italize(RangerJwtAuthHandler.java:84)
~[ranger-authn-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
    at org.apache.ranger.rms.security.RMSJwtAuthFilter.initialize(R
MSJwtAuthFilter.java:77)
~[ranger-rms-common-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
    at jdk.internal.reflect.NativeMethodAccessorImpl.invoke0(Native
Method) ~[?:?]
    at jdk.internal.reflect.NativeMethodAccessorImpl.invoke(NativeMe
thodAccessorImp

```

This exception is logged, since mandatory JWT auth filter configs are not provided (expected for some environments like y-cloud) Even though JWT auth filter is failed to initialised, it falls back to kerberos auth, hence no impact from auth perspective

NA

CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup

This is a warning log printed in catalina.out file when Ranger RMS server is initialised. Following exception is observed only in Oracle 19 setup. FileNotFoundException: /usr/share/java/oraclepki.jar

NA

CDPD-55107: Not able to search using multiple user filter in access audit tab

If you were using multiple user search filters in Audit > Access Tab on Ranger Admin UI, after upgrading to CDP-7.1.9 that would not be supported. You can continue to search user with single search filter.

None

CDPD-48975: Ranger KMS KTS to KMS DB migration : keys with the same name but different case are not migrated

KMS keys are not case sensitive

No work arounds. Such keys combination are very rare and migration doc was updated to check such keys before starting the migration.

CDPD-58704: hadoop roll key / key delete command shows operation failed error when one KMS host is down, even when operation succeeds

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered kms instances. if one KMS instance is stopped (not deleted), client gets runtime exception.

This simply returns the runtime exception on client end for stopped instances but doesn't break any functionality.

CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload:

```
{"resourceName": "classification", "userInput": "", "resources":  {"classification": []}}
```

expectation is to return all the classifications . But the response has only "classification" Happens similarly for entity-label , entity-business-metadata.

None.

CDPD-42598: Kafka policy creation allowed with incorrect permissions.

When creating a Kafka policy from the UI, the permissions "Idempotent write" and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource

None.

CDPD-40734: User allowed to insert data into a hive table when there is a deny policy on a table column.

A user is allowed to enter data into a table even if there is a deny policy present on one of the table columns.

Test scenario details:

```
Policy setup :-
policy 1 :- all access policy for hrt_qa, hive and impala users
resources - database - * , table - * , column - *
users : hrt_qa, hive, impala
access - all access allowed
policy 2 :- policy on test_1.table_1 for hrt_5
users : hrt_5
resources : database - test_1, table - table_1, column - *
access :- all access allowed
policy 3 :- deny policy on test_1.table_1.c0 for hrt_5
users : hrt_5
resources : database - test_1, table - table_1, column - c0
access - all access denied
data setup :-
```

```
database - test_1
table - table_1(c0 int, c1 int)
```

The user is able to insert data into the table.

None.

CDPD-58860: After upgrading CDP-7.1.8 to CDP-7.1.9, cdp-proxy token missing from Knox Ranger policy.

As part of [OPSAPS-67480](#) in 719 default ranger policy is added from cdp-proxy-token topology, so that after a new installation of CDP-7.1.9, the Knox Ranger policy includes cdp-proxy-token. However, upgrades do not add cdp-proxy-token to cm_knox policies automatically.

Manually add cdp-proxy-token to the Knox policy, using Ranger Admin Web UI.

1. Log in to Cloudera Manager Ranger Admin Web UI, as a Ranger administrator.
2. On Ranger Admin Web UI Service Manager Resource Knox, click cm_knox.
3. In Knox Policies, open the CDP Proxy UI, API and Token policy.
4. In Knox Topology*, add cdp-proxy-token.
5. Click Save.
6. Restart Ranger.

CDPD-68806: The Revoke operation for users belonging to a group or role permission does not function as expected.

List command is listing all the tables even when the user permission is revoked. And also the command does not add any deny policy to Ranger for that specific user.

For example:

user_1 is the part of group_1

user_2 is the part of role_1

Create a Ranger policy for group_1 and role_1 to access all the resources in HBase.

Try to revoke the permission for user_1 and user_2 using HBase shell global revoke command:

```
revoke 'user_1'
```

```
revoke 'user_2'
```

But there is no deny policy added for that user. And access is still granted to user_1 and user_2.

This behavior is currently not supported in HBase shell. Must be handled manually using the Ranger policy change.

CDPD-64939: Impala behavior differences across 717 SP3 and 719 SP1 releases.

Upgrading from Cloudera Runtime 7.1.7 Service Pack 3 to 7.1.9 Service Pack 1, for a GRANT statement involving multiple columns in a table, Impala service creates one Ranger policy for each column.

None

CDPD-68739: The revoke command does not work when using the HBase shell

While using the HBase shell, running the revoke command does not cancel the user permission. Users are able to perform actions even after running the revoke command.

For example:

When a user is granted permission to a particular table using the HBase grant command.

```
hbase:003:0> grant 'user_1', 'RWCA', 'some_abc_table'
```


And then execute the global revoke command for that user:

```
hbase:005:0> revoke 'user_1'
```

Permission for that user 'user_1' for the table 'some_abc_table' not got revoked.

None

CDPD-67238: Multiple Columns Revoke not generating policies with correct number of columns

As an example, when "revoke select(col1, col2,col3) on table demo.test from role Role3;" is done, the generated policy does not revoke the columns. Currently the revoke statement is only revoking if there is only one column.

None

CDPD-69412: revokeAccess() behaves differently from secureRevokeAccess() after RANGER-4638

In a non-Kerberized environment, you must use within revokeAccess(), "!isPolicyResourceSameAsRevokedResource" as the condition as it is used in secureRevokeAccess().

None

Known Issues in Ranger KMS

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

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

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

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

None.

Known Issues in Schema Registry

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

CDPD-40380: Authorization checking issue when Kerberos is disabled

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the ranger.plugin.schema-registry.service.name field.
3. Replace GENERATED_RANGER_SERVICE_NAME with the actual name of the service.
4. Restart the Schema Registry service.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

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

Schema Registry is unable to create entities in Atlas if Schema Registry uses an Oracle database. The following will be present in the Schema Registry log if you are affected by this issue:

```
ERROR com.cloudera.dim.atlas.events.AtlasEventsProcessor: An error occurred while processing Atlas events.
java.lang.IllegalArgumentException: Cannot invoke com.hortonworks.registries.schemaregistry.AtlasEventStorable.setType on bean class 'class com.hortonworks.registries.schemaregistry.AtlasEventStorable' - argument type mismatch - had objects of type "java.lang.Long" but expected signature "java.lang.Integer"
```

This issue causes the loss of audit data on Oracle environments.

None.

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

Schema Registry uses a schema fingerprinting mechanism to differentiate between schemas. However, fingerprinting does not take into consideration the schema attributes of the field type. As a result, if you have two schemas where the only difference is that one has type attributes defined and the other does not, they will be considered identical by Schema Registry. For example, the following schemas are considered identical:

```
#Schema V1
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"local_timestamp_micros_long","type":"long"}]}

#Schema V2
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"local_timestamp_micros_long","type":{"type":"long","logicalType":"local-timestamp-micros"}}]}
```

Notice that the only difference is that in the second schema, the `local_timestamp_micros_long` field has a logical type specified. In cases like this, the new version of the schema is not created, the initial version is used. This is true even if the data that is being produced has a new schema version. The ID of the first schema version is used and is put in the serialized record. The new schema version is not created.

This issue is common when using change data capture (CDC) connectors like the Debezium connectors. This is because CDC connectors create schemas with the logical type decimal based on the column type in the database schema. For example:

```
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"database_column","type":{"type":"bytes","logicalType":"decimal","precision":64,"scale":0}}]}
```

If the database schema changes (for example, the column type), it is possible that only scale changes, which is a schema attribute.

```
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"database_column","type":{"type":"bytes","logicalType":"decimal","precision":64,"scale":1}}]}
```

In this case, even though scale changed to 1, the first version of the schema is used where scale is 0. As a result, the data is consumed with the wrong scale.

Avoid using logical types or other attributes. Alternatively, ensure that there are no changes in the logical types or other attributes between schema versions.

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

Schema Registry does not start if the address specified in the Load Balancer Address Ranger property does not end with a trailing slash (/).

Set the value of the `RANGER_REST_URL` Schema Registry environment variable to an address that includes a trailing slash.

1. In Cloudera Manager, select the Schema Registry service.
2. Go to Configuration.
3. Find the Schema Registry Server Environment Advanced Configuration Snippet (Safety Valve) property and add the following:

```
Key: RANGER_REST_URL
```

```
Value: [***RANGER REST API URL***]
```

Replace [***RANGER REST API URL***] with an address that can be used by Schema Registry to access Ranger. Ensure that the address ends with a trailing slash. For example: `https://ranger-1.cloudera.com:6182/`

4. Restart the Schema Registry service.

OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

1. Access the Ranger Console (Ranger Admin web UI).
2. Click the `cm_atlas` resource-based service.
3. Add the `schemaregistry` user to the `all - *` policies.
4. Click `Manage Service Edit Service`.
5. Add the `schemaregistry` user to the `default.policy.users` property.

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

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to `required`. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to `requested` instead of `required` and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Known Issues in Search Client

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

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.

Known Issues in Solr

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

HBase Indexer does not work with JDK 17

Depending on the Cloudera Manager version used with CDP, HBase Indexer (KS Indexer) may have compatibility issues with JDK 17.

You have the following options to fix this issue:

- Upgrade Cloudera Manager to version 7.11.3 or higher.
- If upgrading Cloudera Manager is not an option, you can manually add the following to HBase Indexer Java options in Cloudera Manager:

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

Splitshard operation fails after CDH 6 to CDP upgrade

Collections are not reindexed during an upgrade from CDH 6 to CDP 7 because Lucene 8 (CDP) can read Lucene 7 (CDH 6) indexes.

If you try to execute a SPLITSHARD operation against such a collection, it fails with a similar error message:

```
o.a.s.h.a.SplitOp ERROR executing split: => java.lang.IllegalArg
umentException: Cannot merge a segment t
hat has been created with major version 7 into this index which
has been created by major version 8
    at org.apache.lucene.index.IndexWriter.validateMergeRea
der(IndexWriter.java:3044)
java.lang.IllegalArgumentException: Cannot merge a segment that h
as been created with major version 7 into this index which has b
een created by major version 8
    at org.apache.lucene.index.IndexWriter.validateMergeReade
r(IndexWriter.java:3044) ~[lucene-core-8.11.2.7.1.9.3-2.jar:8.11
.2.7.1.9.3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins -
2023-12-02 00:05:23]
    at org.apache.lucene.index.IndexWriter.addIndexes(IndexWr
iter.java:3110) ~[lucene-core-8.11.2.7.1.9.3-2.jar:8.11.2.7.1.9.
3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins - 2023-12
-02 00:05:23]
    at org.apache.solr.update.SolrIndexSplitter.doSplit(So
lrIndexSplitter.java:318) ~[solr-core-8.11.2.7.1.9.3-2.jar:8.11.
2.7.1.9.3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins -
2023-12-02 00:16:28]
    at org.apache.solr.update.SolrIndexSplitter.split(Solr
IndexSplitter.java:184) ~[solr-core-8.11.2.7.1.9.3-2.jar:8.11.2.
7.1.9.3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins - 2
023-12-02 00:16:28]
    at org.apache.solr.update.DirectUpdateHandler2.split(Dir
ectUpdateHandler2.java:922) ~[solr-core-8.11.2.7.1.9.3-2.jar:8.1
1.2.7.1.9.3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins
- 2023-12-02 00:16:28]
    at org.apache.solr.handler.admin.SplitOp.execute(SplitOp
.java:165) ~[solr-core-8.11.2.7.1.9.3-2.jar:8.11.2.7.1.9.3-2 a6f
f93f9665115dffbdad0ad7f222fd1978d495d - jenkins - 2023-12-02 00:
16:28]
    at org.apache.solr.handler.admin.CoreAdminOperation.execu
te(CoreAdminOperation.java:367) ~[solr-core-8.11.2.7.1.9.3-2.jar
:8.11.2.7.1.9.3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jen
kins - 2023-12-02 00:16:28]
```

This happens because the segment created using a Lucene 7 index cannot be merged into a Lucene 8 index.

Drop the entire collection, delete the data in HDFS and recreate the collection with Solr 8 configs.

Changing the default value of Client Connection Registry HBase configuration parameter causes HBase MRIT job to fail

If the value of the HBase configuration property Client Connection Registry is changed from the default ZooKeeper Quorum to Master Registry then the Yarn job started by HBase MRIT fails with a similar error message:

```
Caused by: org.apache.hadoop.hbase.exceptions.MasterRegistryFetc
hException: Exception making rpc to masters [quasar-bmyccr-2.qua
sar-bmyccr.root.hwx.site,22001,-1]
    at org.apache.hadoop.hbase.client.MasterRegistry.lambda$g
roupCall$1(MasterRegistry.java:244)
    at org.apache.hadoop.hbase.util.FutureUtils.lambda$addLi
stener$0(FutureUtils.java:68)
    at java.util.concurrent.CompletableFuture.uniWhenComple
te(CompletableFuture.java:774)
```

```

        at java.util.concurrent.CompletableFuture.uniWhenComple
eStage(CompletableFuture.java:792)
        at java.util.concurrent.CompletableFuture.whenComplete(Co
mpletableFuture.java:2153)
        at org.apache.hadoop.hbase.util.FutureUtils.addListener(F
utureUtils.java:61)
        at org.apache.hadoop.hbase.client.MasterRegistry.groupCa
ll(MasterRegistry.java:228)
        at org.apache.hadoop.hbase.client.MasterRegistry.call(Ma
sterRegistry.java:265)
        at org.apache.hadoop.hbase.client.MasterRegistry.getMetaR
egionLocations(MasterRegistry.java:282)
        at org.apache.hadoop.hbase.client.ConnectionImplementati
on.locateMeta(ConnectionImplementation.java:900)
        at org.apache.hadoop.hbase.client.ConnectionImplementat
ion.locateRegion(ConnectionImplementation.java:867)
        at org.apache.hadoop.hbase.client.ConnectionImplementati
on.relocateRegion(ConnectionImplementation.java:850)
        at org.apache.hadoop.hbase.client.ConnectionImplementat
ion.locateRegionInMeta(ConnectionImplementation.java:981)
        at org.apache.hadoop.hbase.client.ConnectionImplementa
tion.locateRegion(ConnectionImplementation.java:870)
        at org.apache.hadoop.hbase.client.RpcRetryingCallerWith
ReadReplicas.getRegionLocations(RpcRetryingCallerWithReadReplica
s.java:319)
        ... 21 more
Caused by: org.apache.hadoop.hbase.client.RetriesExhaustedExcept
ion: Failed contacting masters after 1 attempts.
Exceptions:
java.io.IOException: Call to address=quasar-bmyccr-2.quasar-bmy
ccr.root.hwx.site/172.27.19.4:22001 failed on local exception: j
ava.io.IOException: java.lang.RuntimeException: Found no valid a
uthentication method from options
        at org.apache.hadoop.hbase.client.MasterRegistry.lambda
$groupCall$1(MasterRegistry.java:243)
        ... 35 more

```

Add the following line to the MRIT command line:

```
-D 'hbase.client.registry.impl=org.apache.hadoop.hbase.client.ZK
ConnectionRegistry'
```

Solr does not support rolling upgrade to release 7.1.9 or lower

Solr supports rolling upgrades from release 7.1.9 and higher. Upgrading from a lower version means that all the Solr Server instances are shut down, parcels upgraded and activated and then the Solr Servers are started again. This causes a service interruption of several minutes, the actual value depending on cluster size.

Services like Atlas and Ranger that depend on Solr, may face issues because of this service interruption.

None.

Unable to see single valued and multivalued empty string values when querying collections after upgrade to CDP

After upgrading from CDH or HDP to CDP, you are not able to see single valued and multi Valued empty string values in CDP.

This behavior in CDP is due to the remove-blank processor present in solrconfig.xml in Solr 8.

Remove the remove-blank processor from solrconfig.xml.

Cannot create multiple heap dump files because of file name error

Heap dump generation fails with a similar error message:

```
java.lang.OutOfMemoryError: Java heap space
Dumping heap to /data/tmp/solr_solr-SOLR_SERVER-fc9dacc265fabfc5
00b92112712505e3_pid{{PID}}.hprof ...
Unable to create /data/tmp/solr_solr-SOLR_SERVER-fc9dacc265fab
fc500b92112712505e3_pid{{PID}}.hprof: File exists
```

The cause of the problem is that {{PID}} does not get substituted during dump file creation with an actual process ID and because of that, a generic file name is generated. This causes the next dump file creation to fail, as the existing file with the same name cannot be overwritten.

You need to manually delete the existing dump file.

Solr coreAdmin status throws Null Pointer Exception

You get a Null Pointer Exception with a similar stacktrace:

```
Caused by: java.lang.NullPointerException
    at org.apache.solr.core.SolrCore.getInstancePath(SolrCore.
java:333)
    at org.apache.solr.handler.admin.CoreAdminOperation.getCor
eStatus(CoreAdminOperation.java:324)
    at org.apache.solr.handler.admin.StatusOp.execute(StatusOp.
java:46)
    at org.apache.solr.handler.admin.CoreAdminOperation.execute
(CoreAdminOperation.java:362)
```

This is caused by an error in handling solr admin core STATUS after collections are rebuilt.

Restart the Solr server.

Applications fail because of mixed authentication methods within dependency chain of services

Using different types of authentication methods within a dependency chain, for example, configuring your indexer tool to authenticate using Kerberos and configuring your Solr Server to use LDAP for authentication may cause your application to time out and eventually fail.

Make sure that all services in a dependency chain use the same type of authentication.

API calls fail with error when used with alias, but work with collection name

API calls fail with a similar error message when used with an alias, but they work when made using the collection name:

```
[ ] o.a.h.s.t.d.w.DelegationTokenAuthenticationFilter Authenti
cation exception: User: xyz@something.example.com is not allowed
to impersonate xyz@something.example.com
[c:RTOTagMetaOdd s:shard3 r:core_node11 x:RTOTagMetaOdd_shar
d3_replica_n8] o.a.h.s.t.d.w.DelegationTokenAuthenticationFilter
Authentication exception: User: xyz@something.example.com is not
allowed to impersonate xyz@something.example.com
```

Make sure there is a replica of the collection on every host.

CrunchIndexerTool does not work out of the box if /tmp is mounted noexec mode

When you try to run CrunchIndexerTool with the /tmp directory mounted in noexec mode, It throws a snappy-related error.

Create a separate directory for snappy temp files which is mounted with EXEC privileges and set this directory as the value of the org.xerial.snappy.tmpdir java property as a driver java option.

For example:

```
export myDriverJarDir=/opt/cloudera/parcels/CDH//lib/solr/contri
b/crunch;export myDependencyJarDir=/opt/cloudera/parcels/CDH//
```

```
lib/search/lib/search-crunch;export myDriverJar=$(find $myDriverJarDir -maxdepth 1 -name 'search-crunch-*.jar' ! -name '*-job.jar' ! -name '*-sources.jar');export myDependencyJarFiles=$(find $myDependencyJarDir -name '*.jar' | sort | tr '\n' ',' | head -c -1);export myDependencyJarPaths=$(find $myDependencyJarDir -name '*.jar' | sort | tr '\n' ':' | head -c -1);export HADOOP_CONF_DIR=;spark-submit --master local --deploy-mode client --driver-library-path /opt/cloudera/parcels/CDH//lib/hadoop/lib/native/ --jars $myDependencyJarFiles --driver-java-options ' -Dorg.xerial.snappy.tmpdir=/home/systest/tmp ' --class org.apache.solr.crunch.CrunchIndexerTool $myDriverJar --input-file-format=avroParquet --input-file-reader-schema search-parquetfile/parquet-schema.avsc --morphline-file /tmp/mrTestBase.conf --pipeline-type spark --chatty hdfs://[***HOSTNAME***]:8020/tmp/parquetfileparsertest-input
```

Mergeindex operation with --go-live fails after CDH 6 to CDP upgrade

During an upgrade from CDH6 to CDP, collections are not reindexed because Lucene 8 (CDP) can read Lucene 7 (CDH6) indexes.

If you try to execute MapReduceIndexerTool (MRIT) or HBase Indexer MRIT with --go-live against such a collection, you get a similar error message:

```
Caused by: java.lang.IllegalArgumentException: Cannot merge a segment that has been created with major version 8 into this index which has been created by major version 7
    at org.apache.lucene.index.IndexWriter.validateMergeReader(IndexWriter.java:2894)
    at org.apache.lucene.index.IndexWriter.addIndexes(IndexWriter.java:2960)
    at org.apache.solr.update.DirectUpdateHandler2.mergeIndexes(DirectUpdateHandler2.java:570)
    at org.apache.solr.update.processor.RunUpdateProcessor.processMergeIndexes(RunUpdateProcessorFactory.java:95)
    at org.apache.solr.update.processor.UpdateRequestProcessor.processMergeIndexes(UpdateRequestProcessor.java:63)
```

This happens because CDP MRIT and HBase indexer use Solr 8 as embedded Solr, which creates a Lucene 8 index. It cannot be merged (using MERGEINDEXES) into an older Lucene 7 index.

In the case of MRIT the only way to move past this issue is to drop the entire collection, delete the data in HDFS and recreate the collection with Solr 8 configs.

For HBase Indexer MRIT an alternative workaround is setting the number of reducers to 0 (--reducers 0) because in this case documents are sent directly from the mapper tasks to live Solr servers instead of using MERGEINDEXES.

Apache Tika upgrade may break morphlines indexing

The upgrade of Apache Tika from 1.27 to 2.3.0 brought potentially breaking changes for morphlines indexing. Duplicate/triplicate keys names were removed and certain parser class names were changed (For example, org.apache.tika.parser.jpeg.JpegParser changed to org.apache.tika.parser.image.JpegParser).

To avoid morphline commands failing after the upgrade, do the following:

- Check if key name changes affect your morphlines. For more information, see *Removed duplicate/triplicate keys* in [Migrating to Tika 2.0.0](#).
- Check if the name of any parser you use has changed. For more information, see the Apache Tika [API documentation](#).

Update your morphlines if necessary.

CDPD-28006: Solr access via Knox fails with impersonation error though auth_to_local and proxy user configs are set

Currently the names of system users which are impersonating users with Solr should match with the names of their respective Kerberos principals.

If, for some reason, this is not feasible, you must add the user name you want to associate with the custom Kerberos principal to Solr configuration via the Solr Service Environment Advanced Configuration Snippet (Safety Valve) environment variable in Cloudera Manager.

For more information, see [Configuring custom Kerberos principals and custom system users](#).

CDH-77598: Indexing fails with socketTimeout

Starting from CDH 6.0, the HTTP client library used by Solr has a default socket timeout of 10 minutes. Because of this, if a single request sent from an indexer executor to Solr takes more than 10 minutes to be serviced, the indexing process fails with a timeout error.

This timeout has been raised to 24 hours. Nevertheless, there still may be use cases where even this extended timeout period proves insufficient.

If your MapreduceIndexerTool or HBaseMapreduceIndexerTool batch indexing jobs fail with a timeout error during the go-live (Live merge, MERGEINDEXES) phase (This means the merge takes longer than 24 hours).

Use the --go-live-timeout option where the timeout can be specified in milliseconds.

CDPD-12450: CrunchIndexerTool Indexing fails with socketTimeout

The http client library uses a socket timeout of 10 minutes. The Spark Crunch Indexer does not override this value, and in case a single batch takes more than 10 minutes, the entire indexing job fails. This can happen especially if the morphlines contain DeleteByQuery requests.

Try the following workarounds:

- Check the batch size of your indexing job. Sending too large batches to Solr might increase the time needed on the Solr server to process the incoming batch.
- If your indexing job uses deleteByQuery requests, consider using deleteById wherever possible as deleteByQuery involves a complex locking mechanism on the Solr side which makes processing the requests slower.
- Check the number of executors for your Spark Crunch Indexer job. Too many executors can overload the Solr service. You can configure the number of executors by using the --mappers parameter
- Check that your Solr installation is correctly sized to accommodate the indexing load, making sure that the number of Solr servers and the number of shards in your target collection are adequate.
- The socket timeout for the connection can be configured in the morphline file. Add the solrClientSocketTimeout parameter to the solrLocator command

Example

```
SOLR_LOCATOR :
{
  collection : test_collection
  zkHost : "zookeeper1.example.corp:2181/solr"
  # 10 minutes in milliseconds
  solrClientSocketTimeout: 600000
  # Max number of documents to pass per RPC from morphline to
  Solr Server
  # batchSize : 10000
}
```

CDPD-29289: HBaseMapReduceIndexerTool fails with socketTimeout

The http client library uses a socket timeout of 10 minutes. The HBase Indexer does not override this value, and in case a single batch takes more than 10 minutes, the entire indexing job fails.

You can overwrite the default 600000 millisecond (10 minute) socket timeout in HBase indexer using the `--solr-client-socket-timeout` optional argument for the direct writing mode (when the value of the `--reducers` optional argument is set to 0 and mappers directly send the data to the live Solr).

CDPD-20577: Splitshard operation on HDFS index checks local filesystem and fails

When performing a shard split on an index that is stored on HDFS, `SplitShardCmd` still evaluates free disk space on the local file system of the server where Solr is installed. This may cause the command to fail, perceiving that there is no adequate disk space to perform the shard split.

Run the following command to skip the check for sufficient disk space altogether:

- On nonsecure clusters:

```
curl 'http://[${**SOLR_SERVER_HOSTNAME**}]:8983/solr/admin/collections?action=SPLITSHARD&collection=[**COLLECTION_NAME**]&shard=[**SHARD_TO_SPLIT**]&skipFreeSpaceCheck=true'
```

- On secure clusters:

```
curl -k -u : --negotiate 'http://[${**SOLR_SERVER_HOSTNAME**}]:8985/solr/admin/collections?action=SPLITSHARD&collection=[**COLLECTION_NAME**]&shard=[**SHARD_TO_SPLIT**]&skipFreeSpaceCheck=true'
```

Replace `[**SOLR_SERVER_HOSTNAME**]` with a valid Solr server hostname, `[**COLLECTION_NAME**]` with the collection name, and `[**SHARD_TO_SPLIT**]` with the ID of the to split.

To verify that the command executed successfully, check overseer logs for a similar entry:

```
2021-02-02 12:43:23.743 INFO (OverseerThreadFactory-9-thread-5-processing-n:myhost.example.com:8983_solr) [c:example s:shard1] o.a.s.c.a.c.SplitShardCmd Skipping check for sufficient disk space
```

DOCS-5717: 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.

None, you need to reindex collections.

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.

CDH-19407: The Browse and Spell Request Handlers are not enabled in schemaless mode

The Browse and Spell Request Handlers require certain fields to be present in the schema. Since those fields cannot be guaranteed to exist in a Schemaless setup, the Browse and Spell Request Handlers are not enabled by default.

If you require the Browse and Spell Request Handlers, add them to the solrconfig.xml configuration file. Generate a non-schemaless configuration to see the usual settings and modify the required fields to fit your schema.

CDH-17978: 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 indexes. Blockcache writing is disabled by default.

None

CDH-58276: Users with insufficient Solr permissions may receive a "Page Loading" message from the 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 receive a web page that never finishes loading.

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.

Unsupported features

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

- Panel with security info in admin UI's dashboard
- Incremental backup mode
- Schema Designer UI
- Package Management System
- HTTP/2
- Solr SQL/JDBC
- Graph Traversal
- Cross Data Center Replication (CDCR)
- SolrCloud Autoscaling
- HDFS Federation
- Saving search results
- Solr contrib modules (Spark, MapReduce and Lily HBase indexers are not contrib modules but part of the Cloudera Search product itself, therefore they are supported).

Limitations

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.

Known Issues in Spark

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

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

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

The following services support DDL statements:

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

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

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

CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the RuntimeWarning: Failed to add file message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

CDPD-67517: Spark3 tests are failing if /tmp is mounted as NOEXEC.

Map the tmpdir to a writable path in spark3-conf/spark-defaults.conf using the following steps:

1. Go to the Cloudera Manager.
2. Go to the Spark service.
3. Click the Configuration tab.
4. Select Scope Gateway .
5. Select Category Advanced .
6. Locate the Spark Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-defaults.conf_client_config_safety_valve property.
7. Map the tmpdir to a writable path:

```
spark.driver.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
spark.executor.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
```

8. Enter a Reason for change, and then click Save Changes to commit the changes.
9. Deploy the client configuration.

CDPD-23817: In the upgraded Cluster, the permission of /tmp/spark is restricted due to the HDP configuration hive.exec.scratchdir=/tmp/spark.

If you are using the /tmp/spark directory in the CDP cluster, you must provide the required additional Policies/ACL permissions.

CDPD-22670 and CDPD-23103: There are two configurations in Spark, "Atlas dependency" and "spark_lineage_enabled", which are conflicted. The issue is when Atlas dependency is turned off but spark_lineage_enabled is turned on.

Run Spark application, Spark will log some error message and cannot continue. That can be restored by correcting the configurations and restarting Spark component with distributing client configurations.

CDPD-23007: Mismatch in the Spark Default DB Location. In HDP 3.1.5, hive_db entities have one attribute - 'location' which is configured to the '/managed' path. In fresh install of CDP 7.1.7, hive_db entities now have 2 attributes 'location' configured to '/external' path and 'managedLocation' configured to '/managed' path. In the AM2CM migration (HDP 3.1.5 -> CDP 7.1.7), the 'location' attribute from hive_db entities in HDP 3.1.5 comes unaltered to CDP 7.1.7 and hence maps to '/managed' path.

This issue arises only if you are upgrading from HDP 3.1.5 to CDP 7.1.7. If you are performing a fresh install of CDP 7.1.7, you can ignore this issue.

None

CDPD-217: The Apache Spark connector is not supported

The old *Apache Spark - Apache HBase Connector* (shc) is not supported in CDP releases.

Use the new HBase-Spark connector shipped in CDP release.

CDPD-3038: Launching pyspark displays several HiveConf warning messages

When pyspark starts, several Hive configuration warning messages are displayed, similar to the following:

```
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of name hive.meta
store.runworker.in does not exist
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of
name hive.masking.algo does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of
name hive.metastore.runworker.in does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of na
me hive.masking.algo does not exist
```

These errors can be safely ignored.

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

CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

CDPD-11079: Blacklisted topics appear in the list of replicated topics

If a topic was originally replicated but was later excluded for replication, 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, this topic will be included in the response. Additionally, the excluded topic will also be visible in the SMM UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

CDPD-30275: SRM may automatically re-create deleted topics on target clusters

If `auto.create.topics.enable` is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with `srm-control`. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1]
```

2. Wait until SRM is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

Limitations

SRM cannot replicate Ranger authorization policies to or from Kafka clusters

Due to a limitation in the Kafka-Ranger plugin, SRM cannot replicate Ranger policies to or from clusters that are configured to use Ranger for authorization. If you are using SRM to replicate data to or from a cluster that uses Ranger, disable authorization policy synchronization in SRM. This can be achieved by clearing the Sync Topic Acls Enabled (`sync.topic.acls.enabled`) checkbox.

SRM cannot ensure the exactly-once semantics of transactional source topics

SRM data replication uses at-least-once guarantees, and as a result cannot ensure the exactly-once semantics (EOS) of transactional topics in the backup/target cluster.



Note: Even though EOS is not guaranteed, you can still replicate the data of a transactional source, but you must set `isolation.level` to `read_committed` for SRM's internal consumers. This can be done by adding `[***SOURCE CLUSTER ALIAS***]->[***TARGET CLUSTER ALIAS***].consumer.isolation.level=read_committed` to the Streams Replication Manager's Replication Configs SRM service property in Cloudera Manager. The `isolation.level` property can be set on a global connector or replication level. For example:

```
#Global connector level
connectors.consumer.isolation.level
=read_committed
#Replication level
uswest->useast.consumer.isolation.l
evel=read_committed
```

SRM checkpointing is not supported for transactional source topics

SRM does not correctly translate checkpoints (committed consumer group offsets) for transactional topics. Checkpointing assumes that the offset mapping function is always increasing, but with transactional source topics this is violated. Transactional topics have control messages in them, which take up an offset in the log, but they are never returned on the consumer API. This causes the mappings to decrease, causing issues in the checkpointing feature. As a result of this limitation, failover operations for transactional topics is not possible.

Known Issues for Sqoop

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

CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version

If you have upgraded to CDP PvC Base 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to CDP PvC Base 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see [Troubleshooting Apache Sqoop issues](#).

CDPD-44431: Using direct mode causes problems

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the `--direct` option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the `sqoop.enable.deprecated.direct` property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through `-Dsqoop.enable.deprecated.direct=true`.

CDPD-3089: Avro, S3, and HCat do not work together properly

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

Parquet columns inadvertently renamed

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.

PARQUET-99: Importing Parquet files might cause out-of-memory (OOM) errors

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.

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.

OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager

Cloudera Manager does not display a Log Files menu for SMM UI role (and SMM UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by SMM UI is not supported by Cloudera Manager.

View the SMM UI logs on the host.

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in SMM safety valve in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Add the following value for bootstrap servers>Save Changes > Restart SMM.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

CDPD-46728: SMM UI shows the consumerGroup instead of the instances on the Profile page's right hand side

On the ConsumerGroupDetail page, SMM UI shows the group instead of its instances on the right hand side table.

None.

CDPD-47836: The FROM OFFSET field of the offset slider in Data Explorer does not update on partition change

When changing the partition on the **Data Explorer** tab of a topic, the **FROM OFFSET** field of the offset slider is not updated to reflect the first offset of the newly selected partition. If the first offset of the newly selected partition differs from the previous one, the application throws an error. This issue only affects the **Data Explorer** tab, which is accessed from the topic details page of a topic.

This issue does not affect the **Data Explorer** modal window, which is accessed by clicking  next to the name of a topic.

After selecting a new partition, refresh the page. Alternatively, use the **Data Explorer** modal window.

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

KNOX does not send neither the Authentication header nor the hadoop.auth cookie. Because of this, the SMM UI returns an HTTP 401 response and sets the www-authenticate: Negotiate header. After this, KNOX still does not send the Authentication header. This results in the SMM UI being inaccessible through Knox.

You can access the SMM UI directly in Cloudera Manager at Clusters SMM Streams Messaging Manager Web UI and log in using the Kerberos username and password.

CDPD-82560: Streams Messaging Manager known unprotected endpoint

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

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

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 MapReduce, YARN and YARN Queue Manager

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

Known Issues**COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."**

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

None.

COMPX-13177: QueueManager webapp requests fail with 'HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA'

Products:

- Cloudera Manager for CDP Private
- Cloud Base Cloudera Manager for CDP Public Cloud

Context:

- Centos 7.8 and Redhat 7.8 operating systems, when FIPS support is enabled.

Problem:

- When attempting to display the Yarn Queue Manager interface, Cloudera Manager displays an error: "HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA".

1. Edit the file [/etc/default/cloudera-scm-server]

2. Around line 28, modify the line that starts with

```
#export CMF_OVERRIDE_TLS_CIPHERS=...
```

3. Remove the comment mark #.
4. Remove all ciphers with "3DES" in the name.
5. Save the file.
6. Restart the Cloudera Manager Server service.

COMPX-4644: Queue capacity rounding problem when configuration is initially set via YARN

When setting the capacity scheduler configuration through the YARN/Cloudera Manager configuration, there may be capacity values that use multiple decimal places. This results in rounding/floating point precision discrepancies in the UI when trying to validate that all sibling capacities equal 100%. The UI looks like all the numbers add up to 100, but the validation still displays an error and does not allow to save the capacities. It is also observed that the capacity is being calculated as, for example, 99.9999999991 in the backend.

- Create queues within the UI, or
- Ensure that capacities configured through the Capacity Scheduler safety valve do not have more than one decimal place.

20202 Database migration after enabling opt-in migration

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to CDP 7.1.9, forcing migration from H2 to PostgreSQL database.
- Upgrade to CDP 7.1.9 CHF2, moving back to H2 database.
- Upgrade to CDP 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version CDP 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

CDPD-56559: MapReduce jobs can intermittently fail during a rolling upgrade.

During a rolling upgrade between CDP versions 7.1.8 and 7.1.9, MapReduce jobs may fail with message, `RuntimeException: native snappy library not available`. Although the native Snappy compression library is not loaded, a checkmark displays to indicate that the Snappy compression library is loading for NodeManagers that are pending upgrades. This causes the MapReduce jobs that are associated with the NodeManagers to fail. After the upgrade, the jobs work as expected. This issue only impacts rolling upgrades from before CDP 7.1.9 to a higher version.

None.

COMPX-12021 Queue Manager configurations on Scheduler Configuration page are not working

When setting the following properties on the YARN Queue Manager UI, the properties are set in the `capacity-scheduler.xml` file which does not have any effect on YARN. The properties need to be set in the `yarn-site.xml` file, which does not happen when you set them through YARN Queue Manager.

- "Maximum Application Priority" – "yarn.cluster.max-application-priority"
- "Enable Monitoring Policies" – "yarn.resourcemanager.scheduler.monitor.enable"
- "Monitoring Policies" – "yarn.resourcemanager.scheduler.monitor.policies"



Note: You can also set this property on the YARN configuration page in Cloudera Manager as "Capacity Scheduler Preemption".

- "Preemption: Observe Only" – "yarn.resourcemanager.monitor.capacity.preemption.observe_only"

- "Preemption: Monitoring Interval (ms)" – `yarn.resourcemanager.monitor.capacity.preemption.monitoring_interval`
 - "Preemption: Maximum Wait Before Kill (ms)" – `yarn.resourcemanager.monitor.capacity.preemption.max_wait_before_kill`
 - "Preemption: Total Resources Per Round" – `yarn.resourcemanager.monitor.capacity.preemption.total_preemption_per_round`
 - "Preemption: Over Capacity Tolerance" – `yarn.resourcemanager.monitor.capacity.preemption.max_ignored_over_capacity`
 - "Preemption: Maximum Termination Factor" – `yarn.resourcemanager.monitor.capacity.preemption.natural_termination_factor`
 - "Enable Intra Queue Preemption" – `yarn.resourcemanager.monitor.capacity.preemption.intra-queue-preemption.enabled`
1. In Cloudera Manager, select the YARN service.
 2. Click the Configuration tab.
 3. Search for `yarn-site.xml`.
 4. Under YARN Service Advanced Configuration Snippet (Safety Valve) for `yarn-site.xml`, add the corresponding parameter and value you need.
 5. Click Save Changes.
 6. Restart the YARN services.

COMPX-11380 YARN Queue Manager displays an error stating that it is unable to complete a request

If YARN Queue Manager incorrectly detects an inactive Resource Manager, it displays an error stating that it is unable to complete a request.

Refresh the browser and reload the YARN Queue Manager UI.

COMPX-6214: When there are more than 600 queues in a cluster, potential timeouts occur due to performance reasons that are visible in the Configuration Service.

The Cloudera Manager proxy timeout configuration field is added now. This issue is tracked in OPSAPS-60554. For this release, the timeout is increased from 20 seconds to 5 minutes. However, if this problem occurs, Cloudera recommends you to increase the proxy timeout value.

COMPX-5817: YARN Queue Manager UI is not able to present a view of pre-upgrade queue structure. Cloudera Manager Store is not supported and therefore YARN does not have any of the pre-upgrade queue structure preserved.

When a Data Hub cluster is deleted, all saved configurations are also deleted. All YARN configurations are saved in Cloudera Manager Store and this is yet to be supported in Data Hub and Cloudera Manager. Hence, the YARN queue structure is also lost when a Data Hub cluster is deleted or upgraded or restored.

COMPX-6628: Unable to delete single leaf queue assigned to a partition.

In the current implementation, you cannot delete a single leaf queue assigned to a partition.

For each non-default partition, perform the following for the single child leaf queue and its parent queues:

1. In Cloudera Manager, click Cluster YARN .
2. Click the **Configuration** tab.
3. Search for ResourceManager. In the Filters pane, under Scope, select ResourceManager.
4. Add the following in Capacity Scheduler Configuration Advanced Configuration Snippet (Safety Valve):

```
Name: yarn.scheduler.capacity.<queuePath>.accessible-node-labels.<partition>.capacity
```

Value: 0

Set the value to 0 in Percentage mode, and 0w in Weight mode, and [memory=0,vcores=0] in Absolute mode.

Name: yarn.scheduler.capacity.<queuePath>.accessible-node-labels.<partition>.maximum-capacity
Value: 100

Set the value to 100 in Percentage and Weight mode and [memory=0,vcores=0] in Absolute mode.

5. Adjust the capacities of the siblings of the parent queue for the same partition.
6. Click Save Changes.
7. Restart the active **ResourceManager** service for the changes to apply.
8. In **Cloudera Manager**, click Cluster YARN Queue Manager UI .
9. Delete the desired single child leaf queue.

COMPX-5240: Restarting parent queue does not restart child queues in weight mode

When a dynamic auto child creation enabled parent queue is stopped in weight mode, its static and dynamically created child queues are also stopped. However, when the dynamic auto child creation enabled parent queue is restarted, its child queues remain stopped. In addition, the dynamically created child queues cannot be restarted manually through the YARN Queue Manager UI either.

Delete the dynamic auto child creation enabled parent queue. This action also deletes all its child queues, both static and dynamically created child queues, including the stopped dynamic queues. Then recreate the parent queue, enable the dynamic auto child creation feature for it and add the required static child queues.

COMPX-5589: Unable to add new queue to leaf queue with partition capacity in Weight/Absolute mode

Scenario

1. You create one or more partitions.
2. Assign a partition to a parent with children
3. Switch to the partition to distribute the capacities
4. Create a new child queue under one of the leaf queues but the following error is displayed:

```
Error :
2021-03-05 17:21:26,734 ERROR
com.cloudera.cpx.server.api.repositories.SchedulerRepository:
Validation failed for Add queue
operation. Error message: CapacityScheduler configuration
validation failed:java.io.IOException:
Failed to re-init queues : Parent queue 'root.test2' have
children queue used mixed of weight
mode, percentage and absolute mode, it is not allowed, please
double check, details:
{Queue=root.test2.test2childNew, label= uses weight mode}.
{Queue=root.test2.test2childNew,
label=partition uses percentage mode}
```

To create new queues under leaf queues successfully, perform the following:

1. Switch to Relative mode
2. Create the required queues
3. Create the required partitions
4. Assign partitions and set capacities
5. Switch back to Weight mode
1. Create the entire queue structure

2. Create the required partitions
3. Assign partition to queues
4. Set partition capacities

COMPX-5264: Unable to switch to Weight mode on creating a managed parent queue in Relative mode

In the current implementation, if there is an existing managed queue in Relative mode, then conversion to Weight mode is not allowed.

To proceed with the conversion from Relative mode to Weight mode, there should not be any managed queues. You must first delete the managed queues before conversion. In Weight mode, a parent queue can be converted into managed parent queue.

COMPX-5549: Queue Manager UI sets maximum-capacity to null when you switch mode with multiple partitions

If you associate a partition with one or more queues and then switch the allocation mode before assigning capacities to the queues, an Operation Failed error is displayed because `max-capacity` is set to null.

After you associate a partition with one or more queues, in the YARN Queue Manager UI, click Overview [***PARTITION NAME***] from the drop-down list and distribute capacity to the queues before switching allocation mode or creating placement rules.

COMPX-4992: Unable to switch to absolute mode after deleting a partition using YARN Queue Manager

If you delete a partition (node label) which has been associated with queues and those queues have capacities configured for that partition (node label), the `CS.xml` still contains the partition (node label) information. Hence, you cannot switch to absolute mode after deleting the partition (node label).

It is recommended not to delete a partition (node label) which has been associated with queues and those queues have capacities configured for that partition (node label).

COMPX-1445: YARN Queue Manager operations are failing when Queue Manager is installed separately from YARN

If Queue Manager is not selected during YARN installation, Queue Manager operations are failing. Queue Manager says 0 queues are configured and several failures are present. That is because ZooKeeper configuration store is not enabled.

1. In Cloudera Manager, select the YARN service.
2. Click the Configuration tab.
3. Find the Queue Manager Service property.
4. Select the Queue Manager service that the YARN service instance depends on.
5. Click Save Changes.
6. Restart all services that are marked stale in Cloudera Manager.

COMPX-1451: Queue Manager does not support multiple ResourceManagers

When YARN High Availability is enabled there are multiple ResourceManagers. Queue Manager receives multiple ResourceManager URLs for a High Availability cluster. It picks the active ResourceManager URL only when Queue Manager page is loaded. Queue Manager cannot handle it gracefully when the currently active ResourceManager goes down while the user is still using the Queue Manager UI.

Reload the Queue Manager page manually.

COMPX-3329: Autorestart is not enabled for Queue Manager in Data Hub

In a Data Hub cluster, Queue Manager is installed with autorestart disabled. Hence, if Queue Manager goes down, It does not restart automatically.

If Queue Manager goes down in a Data Hub cluster, you must go to the Cloudera Manager Dashboard and restart the Queue Manager service.

Third-party applications do not launch if MapReduce framework path is not included in the client configuration

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the `mapreduce.application.framework.path` property is set to the appropriate value, but third-party applications with their own configurations does not launch.

Set the `mapreduce.application.framework.path` property to the appropriate configuration for third-party applications.

JobHistory URL mismatch after server relocation

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated the `mapred-site.xml` file that references the correct JobHistory Server.

CDH-49165: History link in ResourceManager web UI broken for killed Spark applications

When a Spark application is killed, the history link in the ResourceManager web UI does not work.

To view the history for a killed Spark application, see the Spark HistoryServer web UI instead.

CDH-6808: Routable IP address required by ResourceManager

ResourceManager requires routable host:port addresses for `yarn.resourcemanager.scheduler.address`, and does not support using the wildcard `0.0.0.0` address.

Set the address, in the form `host:port`, either in the client-side configuration, or on the command line when you submit the job.

YARN cannot start if Kerberos principal name is changed

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN does not start. In such cases, the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the `znode` and restart the YARN service.
- Use the `reset ZK ACLs` command. This also sets the `znodes` below `/rmstore/ZKRMStateRoot` to `world:anyone:cdw` which is less secure.

Queue Manager does not open on using a custom user with a default Kerberos principal

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager `process_username` property matches the YARN `process_username` property.

COMPX-8687: Missing access check for getAppAttempts

When the Job ACL feature is enabled using Cloudera Manager (YARN Configuration Enable JOB ACL property), the `mapreduce.cluster.acls.enabled` property is not generated to all configuration files, including the `yarn-site.xml` configuration file. As a result, the ResourceManager process uses the default value of this property. The default property of `mapreduce.cluster.acls.enabled` is `false`.

Enable the Job ACL feature using the Advanced Configuration Snippet:

1. In Cloudera Manager select the YARN service.
2. Click Configuration.
3. Find the YARN Service MapReduce Advanced Configuration Snippet (Safety Valve) property.

4. Click the plus icon and add the following:
 - Name: mapreduce.cluster.acls.enabled
 - Value: true
5. Click Save Changes.

COMPX-7493: YARN Tracking URL that is shown in the command line does not work when Knox is enabled

When Knox is configured for YARN, the Tracking URL printed in the command line of an YARN application such as spark-submit shows the direct URL instead of the Knox Gateway URL.

Upgrade CDP Runtime to CDP 7.1.9 CHF 2, and perform the following steps:

1. Open the Cloudera Manager Admin Console and go to the Knox service.
2. Click on the Knox Gateway Home URL.
3. Copy the YARN Resource Manager Web UI V2 URL from the Knox Gateway Home page.

For example, `https://knox-gateway.example.com:8443/gateway/cdp-proxy/yarnuiv2/`

4. Open the Cloudera Manager Admin Console and go to the YARN service.
5. Click the Configuration tab and search for `resourcemanager_config_safety_valve`.
6. Add the Resource Manager Advanced Configuration Snippet (Safety Valve) for `yarn-site.xml` property, and specify its value by using the YARN Resource Manager Web UI V2 URL, copied earlier, as follows:

```
Name: yarn.web-proxy.gateway.url
Value:[***YARN Resource Manager Web UI V2 URL***]
```

7. Enter a Reason for Change and then click Save Changes.
8. Restart YARN.

Unsupported Features

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

- Application Timeline Server v2 (ATSV2)
- Auxiliary Services
- Container Resizing
- Distributed or Centralized Allocation of Opportunistic Containers
- Distributed Scheduling
- Docker on YARN (DockerContainerExecutor)
- 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 Zeppelin

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

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.

BUG-125263: Zeppelin service move fails on clusters upgraded from HDP3.1.5

Resolve the circular symlink issue on the Zeppelin node by linking the conf directory to a new directory under /etc/zeppelin:

- # mkdir -p /etc/zeppelin/<version>/0
- # rm /usr/hdp/<version>/zeppelin/conf
- # ln -s /etc/zeppelin/<version>/0 /usr/hdp/<version>/zeppelin/conf

Where version is the HDP version. For example, 7.1.x-yyy. Restart the Zeppelin server in Ambari.

CDPD-3090: Due to a configuration typo, functionality involving notebook repositories does not work

Due to a missing closing brace, access to the notebook repositories API is blocked by default.

From the CDP Management Console, go to Cloudera Manager for the cluster running Zeppelin. On the Zeppelin configuration page (Zeppelin serviceConfiguration), enter shiro urls in the Search field, and then add the missing closing brace to the notebook-repositories URL, as follows:

```
/api/notebook-repositories/** = authc, roles[{{zeppelin_admin_group}}]
```

Click Save Changes, and restart the Zeppelin service.

CDPD-2406: Logout button does not work

Clicking the Logout button in the Zeppelin UI logs you out, but then immediately logs you back in using SSO.

Close the browser.

OPSAPS-59802: Zeppelin and Livy roles should be co-located on the same host.

When installing or upgrading to CDP Private Cloud Base, you must co-locate all Zeppelin and Livy roles on the same cluster host due to an issue with certificate generation.

Known Issues in ZooKeeper

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

Zookeeper-client does not use ZooKeeper TLS/SSL automatically

The command-line tool 'zookeeper-client' is installed on 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.

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
.zookeeper.ClientCnxnSocketNetty -Dzookeeper.ssl.keyStore
e.location=[***PATH TO YOUR CONFIGURED KEYSTORE***] -Dz
ookeeper.ssl.keyStore.password=[***PASSWORD YOU CONFIGURED
FOR KEYSTORE***] -Dzookeeper.ssl.trustStore.location
n=[***PATH TO YOUR CONFIGURED TRUSTSTORE***] -Dzookeeper
.ssl.trustStore.password=[***THE PASSWORD YOU CONFIGURED
FOR THE TRUSTSTORE***] -Dzookeeper.client.secure=true" zoo
keeper-client -server [***YOUR.ZOOKEEPER.SERVER-1***]:218
2,[***YOUR.ZOOKEEPER.SERVER-2***]:218
2,[***YOUR.ZOOKEEPER.SERVER-3***]:2182
```

TLS v1.3 Support for CDP in Zookeeper

ZooKeeper does not support TLS v1.3.

To configure TLS v1.3, add the following properties in the Server Advanced Configuration Snippet (Safety Valve) for zoo.cfg file.

```
ssl.protocol=TLSv1.3
ssl.enabledProtocols=TLSv1.3,TLSv1.2
ssl.ciphersuites=TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,TLS_E
CDHE_RSA_WITH_AES_128_GCM_SHA256,TLS_ECDHE_ECDSA_WITH_AES_256_GC
M_SHA384,TLS_E
CDHE_RSA_WITH_AES_256_GCM_SHA384,TLS_ECDHE_ECDSA_WITH_CHACHA20_PO
LY1305_SHA256,TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256,TLS_DH
E_RSA_WITH_
AES_128_GCM_SHA256,TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,TLS_AES_1
28_GCM_SHA256,TLS_AES_256_GCM_SHA384,TLS_CHACHA20_POLY1305_SHA256
```

Behavioral changes in Cloudera Runtime 7.1.9 SP1

You can review the changes in certain features or functionalities of components that have resulted in a change in behavior from the previously released version to this version of Cloudera Runtime 7.1.9 SP1.

Behavioral Changes in Ozone

Learn about the change in certain functionality of Ozone that has resulted in a change in behavior from the previously released version to this version of Cloudera Runtime.

Change in the behaviour of the Ozone admin pipeline list command.

In the CDP Private Cloud Base 7.1.9, running the Ozone admin pipeline list command displays the following output: Pipeline[Id: pipeline_id, Nodes: UUID1(hostname1/ip1)UUID2(hostname2/ip2)UUID3(hostname3/ip3), ReplicationConfig: RATIS/THREE, State:OPEN, leaderId:id, CreationTimestamp].

In the CDP Private Cloud Base 7.1.9 SP1, running the Ozone admin pipeline list command displays the following output: Pipeline[Id: pipeline_id, Nodes: UUID1(hostname1/ip1) ReplicaIndex: UUID2(hostname2/ip2) ReplicaIndex: UUID3(hostname3/ip3) ReplicaIndex: 0, ReplicationConfig: RATIS/THREE, State:OPEN, leaderId:id, CreationTimestamp].

CDP Private Cloud Base API Modifications and Removals

All the APIs that have either been modified or removed for this release of CDP Private Cloud Base.

CDP 7.1.9 SP1 and 7.1.9 SP1 Cumulative Hotfix 1 Components with API differences

The following is a list of all the APIs that have been modified or removed in this release of CDP Private Cloud Base. The APIs listed here are on a component by component basis and may include recommendations for how to manage the change.

API Compatibility changes in 7.1.9 SP1 Cumulative Hotfix 1 for HBase

Removed or Modified APIs in CDP 7.1.9 SP1 Cumulative Hotfix 1 for HBase and recommendations for how to handle them.

Apache Version of HBase in 7.1.9 was 1.0.0 and Apache Version of HBase in 7.1.9 SP1 Cumulative Hotfix 1 is 1.0.0

Removed APIs in 7.1.9 SP1 Cumulative Hotfix 1

The following APIs are no longer available for HBase in CDP 7.1.9 SP1 Cumulative Hotfix 1.

OAuthBearerSignedJwt.<init>

Method Removed

Package Name


```
com.cloudera.hbase.security.oauthbearer.internals.knox
```

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

Backport fix for CDPD-67520: JWT authentication expects [sub] claim in the payload. Initially introduced by 7.1.9 CHF6

Recommendation

N/A

Recompilation Required?

Yes

CDP 7.1.9 GA and 7.1.9 SP1 Components with API differences

The following is a list of all the APIs that have been modified or removed in this release of CDP Private Cloud Base. The APIs listed here are on a component by component basis and may include recommendations for how to manage the change.

API Compatibility changes in 7.1.9 SP1 for Hadoop

Removed or Modified APIs in CDP 7.1.9 SP1 for Hadoop and recommendations for how to handle them.

Apache Version of Hadoop in 7.1.9 was 3.1.1 and Apache Version of Hadoop in 7.1.9 SP1 is 3.1.1

Modified APIs in 7.1.9 SP1

The following APIs have been modified for Hadoop and include a description of the impact of the modification on their use.

YarnConfiguration

Field NM_LINUX_CONTAINER_GROUP of type java.lang.String with the compile-time constant value "yarn.nodemanager.linux-container-executor.group" has been removed from this class.

Package Name

```
org.apache.hadoop.yarn.conf
```

Effect

Recompilation of a client program may be terminated with the message: cannot find variable NM_LINUX_CONTAINER_GROUP in org.apache.hadoop.yarn.conf.YarnConfiguration.

Reason for change

Unused constant ([YARN-4599](#))

Recommendation

N/A

Recompilation Required?

Yes

API Compatibility changes in 7.1.9 SP1 for Atlas

Removed or Modified APIs in CDP 7.1.9 SP1 for Atlas and recommendations for how to handle them.

Apache Version of Atlas in 7.1.9 was 3.0.0 and Apache Version of Atlas in 7.1.9 SP1 is 3.0.0

Removed APIs in 7.1.9 SP1

The following APIs are no longer available for Atlas in CDP 7.1.9 SP1

AtlasClientV2.<init>

The method isn't removed, but modified to include a different parameter. Change was done as part of Jersey Upgrade to 2.0, not yet merged in Apache

Package Name

```
org.apache.atlas
```

Effect

The method isn't removed

Reason for change

but modified to include a different parameter. Change was done as part of Jersey Upgrade to 2.0

Recommendation

not yet merged in Apache.

Recompilation Required?

The client application needs to pass WebTarget from jersey 2.0 object as argument instead from WebResource of jersey 1.0

Modified APIs in 7.1.9 SP1

The following APIs have been modified for Atlas and include a description of the impact of the modification on their use.

AdminResource.deleteRule

The method was not removed, rather modified to return the list of deleted rules and the API path was changed because it was conflicting with other API path, which was added in this release (ATLAS-4797)

Package Name

org.apache.atlas.web.resources

Effect

Any client application invoking this API needs to update the API path and if required, they can use the response object

Reason for change

Yes

Recommendation

The method was not removed

Recompilation Required?

rather modified to return the list of deleted rules and the API path was changed because it was conflicting with other API path

CDP 7.1.9 Cumulative hotfix 7 and 7.1.9 SP1 Components with API differences

The following is a list of all the APIs that have been modified or removed in this release of CDP Private Cloud Base. The APIs listed here are on a component by component basis and may include recommendations for how to manage the change.

API Compatibility changes in 7.1.9 SP1 for Hadoop

Removed or Modified APIs in CDP 7.1.9 SP1 for Hadoop and recommendations for how to handle them.

Apache Version of Hadoop in 7.1.9 was 3.1.1 and Apache Version of Hadoop in 7.1.9 SP1 is 3.1.1

Modified APIs in 7.1.9 SP1

The following APIs have been modified for Hadoop and include a description of the impact of the modification on their use.

YarnConfiguration

Field NM_LINUX_CONTAINER_GROUP of type java.lang.String with the compile-time constant value "yarn.nodemanager.linux-container-executor.group" has been removed from this class.

Package Name

org.apache.hadoop.yarn.conf

Effect

Recompilation of a client program may be terminated with the message: cannot find variable NM_LINUX_CONTAINER_GROUP in org.apache.hadoop.yarn.conf.YarnConfiguration.

Reason for change

[YARN-4599](#) removed this unused variable. It seems unrelated to the change itself.

Recommendation

Details about this config option: The UNIX group that the linux-container-executor should run as. This is intended to be set as part of container-executor.cfg. If a YARN client needs this variable it shall be created manually: public static final String NM_LINUX_CONTAINER_GROUP=""yarn.nodemanager.linux-container-executor.group"";

Recompilation Required?

Yes

API Compatibility changes in 7.1.9 SP1 for Atlas

Removed or Modified APIs in CDP 7.1.9 SP1 for Atlas and recommendations for how to handle them.

Apache Version of Atlas in 7.1.9 was 3.0.0 and Apache Version of Atlas in 7.1.9 SP1 is 3.0.0

Removed APIs in 7.1.9 SP1

The following APIs are no longer available for Atlas in CDP 7.1.9 SP1.

AtlasClientV2.<init>

Method Removed

Package Name

org.apache.atlas

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

The method isn't removed, but modified to include a different parameter. The change was done as part of Jersey Upgrade to 2.0, which is not yet merged in Apache.

Recommendation

The client application needs to pass WebTarget from jersey 2.0 object as argument instead of WebResource of jersey 1.0.

Recompilation Required?

Yes

Modified APIs in 7.1.9 SP1

The following APIs have been modified for Atlas and include a description of the impact of the modification on their use.

AdminResource.deleteRule

Return value type has been changed from void to org.apache.atlas.model.instance.EntityMutationResponse.

Package Name

org.apache.atlas.web.resources

Effect

This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

Reason for change

The method was not removed, rather modified to return the list of deleted rules and the API path was changed because it was conflicting with other API path, which was added in this release ([ATLAS-4797](#))

Recommendation

Any client application invoking this API needs to update the API path and if required, they can use the response object.

Recompilation Required?

Yes

CDP 7.1.9 SP1 Components with API differences

The following is a list of all the APIs that have been modified or removed in this release of CDP Private Cloud Base. The APIs listed here are on a component by component basis and may include recommendations for how to manage the change.

API Compatibility changes in 7.1.9 SP1 for ZooKeeper

Removed or Modified APIs in CDP 7.1.9 SP1 for ZooKeeper and recommendations for how to handle them.

Apache Version of ZooKeeper in 7.1.9 SP1 was 3.8.1 and Apache Version of ZooKeeper in 7.1.9 SP1 is 3.8.1

Removed APIs in 7.1.9 SP1

The following APIs are no longer available for ZooKeeper in CDP 7.1.9 SP1

[*SecurityUtils.createSaslClient*](#)

Method Removed

Package Name

org.apache.zookeeper.util

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[ZOOKEEPER-4889](#) : ZooKeeper SecurityUtils has hard coded SASL DIGEST-MD5 - Fallback to DIGEST-MD5 auth mech should be disabled in Fips mode (branch-3.8)

Recommendation

Client code need to be adjusted to supply ZKConfig to the method.

Recompilation Required?

Yes

Deprecation notices in Cloudera Runtime 7.1.9 SP1

Certain features and functionalities have been removed or deprecated in Cloudera Runtime 7.1.9 SP1. You must review these items to understand whether you must modify your existing configuration. You can also learn about the features that will be removed or deprecated in the future release to plan for the required changes.

Terminology

Items in this section are designated as follows:

Deprecated

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

Moving

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

Removed

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

You must contact Cloudera Support or your Cloudera Account Team if you have any questions.

Platform and OS

The listed Operating Systems and databases are deprecated or removed from the Cloudera Runtime 7.1.9 SP1 release.

Database Support:

The following databases are **removed** and no longer supported from the Cloudera Runtime 7.1.9 SP1 release:

- PostgreSQL 11

The following databases are **deprecated** from the Cloudera Runtime 7.1.9 SP1 release:

- PostgreSQL 12
- MariaDB 10.4
- MySQL 5.7

Operating System

The following operating systems are **deprecated** from the Cloudera Runtime 7.1.9 SP1 release:

- RHEL 8.6
- RHEL 7.9
- RHEL 7.9 (FIPS)
- CentOS 7.9
- SLES 12 SP5

Ozone

Certain features and functionality in Ozone are deprecated or removed in Cloudera Runtime 7.1.9 SP1 release.

Removed

Ordered deletion of snapshots

Starting from Cloudera Runtime 7.1.9 SP1 CHF12 onwards, ordered deletion is removed after significant fixes and improvements to Snapshot Garbage Collection. Configurations related to ordered deletion such as `ozone.snapshot.ordered.deletion.enabled` and `ozone.snapshot.deep.cleaning.enabled` might still be found in Cloudera Manager (CM) in order to ensure backward compatibility but they are ignored in versions from 7.1.9 SP1 CHF12 onwards.

Fixed Common Vulnerabilities and Exposures 7.1.9 SP1

Common Vulnerabilities and Exposures (CVE) that is fixed in this release.

- [CVE-2023-39410](#) Apache Avro
- [CVE-2023-46120](#) Amqp-client
- [CVE-2023-36478](#) Eclipse Jetty
- [CVE-2023-26048](#) Eclipse Jetty
- [CVE-2023-26049](#) Eclipse Jetty
- [CVE-2023-40167](#) Eclipse Jetty
- [CVE-2023-41900](#) Eclipse Jetty
- [CVE-2023-38493](#) Armeria
- [CVE-2023-2976](#) Guava
- [CVE-2018-10237](#) Guava
- [CVE-2020-8908](#) Guava
- [CVE-2021-31684](#) json-smart
- [CVE-2023-1370](#) json-smart
- [CVE-2021-27568](#) json-smart
- [CVE-2022-46337](#) Apache Derby

- [CVE-2023-22006](#) Graal-sdk
- [CVE-2023-39196](#) Apache Ozone
- [CVE-2024-23944](#) Apache Zookeeper
- [CVE-2022-46751](#) Apache Ivy
- [CVE-2023-37895](#) Jackrabbit
- [CVE-2021-33813](#) Jdom
- [CVE-2022-40149](#) Jettison
- [CVE-2022-40150](#) Jettison
- [CVE-2022-45685](#) Jettison
- [CVE-2022-45693](#) Jettison
- [CVE-2023-1436](#) Jettison
- [CVE-2022-2048](#) Eclipse Jetty
- [CVE-2022-2047](#) Eclipse Jetty
- [CVE-2023-31582](#) Jose4j
- [CVE-2023-5072](#) org.json
- [CVE-2023-34610](#) json-io
- [CVE-2021-4178](#) kubernetes
- [CVE-2023-34054](#) Reactor-netty
- [CVE-2023-34062](#) Reactor-netty
- [CVE-2023-34478](#) Apache Shiro
- [CVE-2023-46749](#) Apache Shiro
- [CVE-2023-46750](#) Apache Shiro
- [CVE-2023-43642](#) Snappy Java
- [CVE-2023-34034](#) Spring Security
- [CVE-2023-34035](#) Spring Security
- [CVE-2024-22257](#) Spring Security
- [CVE-2023-41080](#) Apache Tomcat
- [CVE-2023-42794](#) Apache Tomcat
- [CVE-2023-42795](#) Apache Tomcat
- [CVE-2023-45648](#) Apache Tomcat
- [CVE-2024-1300](#) vertx
- [CVE-2023-41329](#) Wiremock-jre8
- [CVE-2023-41327](#) Wiremock-jre8
- [CVE-2023-44483](#) xmlsec
- [CVE-2024-45478](#) - Apache Ranger
- [CVE-2024-45479](#) - Apache Ranger

Release Matrix

You must be familiar with the versions of all the cumulative hotfixes for CDP Private Cloud Base version 7.1.9 SP1.

The Runtime Release build number has two parts, the [**Runtime version number**] and the [**Runtime Build number**]. For example, in 7.1.9.1000-103, 7.1.9 is the version number and 103 is the build number.

Table 2: Table 1: Runtime Cumulative hotfixes

Cumulative Hotfix	Runtime Release Build Number	GBN	Release Date	Download URL
7.1.9 SP1 CHF 15	Runtime 7.1.9.1073-8	76451206	Mar 13, 2026	https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1073/parcels/
7.1.9 SP1 CHF 14	Runtime 7.1.9.1069-3	74330271	Jan 08, 2026	https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1069/parcels/
7.1.9 SP1 CHF 13	Runtime 7.1.9.1066-8	73837803	Dec 15, 2025	https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1066/parcels/
7.1.9 SP1 CHF 12	Runtime 7.1.9.1064-1	71379092	Oct 21, 2025	https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1064/parcels/
7.1.9 SP1 CHF 11	Runtime 7.1.9.1059-4	70393529	Sep 10, 2025	https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1059/parcels/
7.1.9 SP1 CHF 10	Runtime 7.1.9.1054-4	69605909	Aug 13, 2025	https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1054/parcels/
7.1.9 SP1 CHF 9	Runtime 7.1.9.1045-5	67903105	Jul 15, 2025	https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1045/parcels/
		407		https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1045/parcels/

Table 3: Cloudera Runtime versions certified with Cloudera Data Services on premises

Cumulative Hotfix	Runtime Release Build Number	GBN	Release Date	Download URL
7.1.9 SP1 CHF11	7.1.9.1059-4	70393529	Sep 10, 2025	https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1059/parcels/
7.1.9 SP1 CHF5	7.1.9.1040-4	65818059	Apr 28, 2025	https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1040/parcels/
7.1.9 SP1 CHF4	7.1.9.1023-3	60818430	Jan 13, 2025	https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1023/parcels/

Cumulative hotfixes

You can review the list of cumulative hotfixes that were shipped for CDP Private Cloud Base version 7.1.9 SP1.

Cumulative hotfix 15

Know more about the cumulative hotfix 15 for CDP 7.1.9 SP1. This cumulative hotfix was released on March 13, 2026.

Cloudera Runtime 7.1.9.1073-8 (Cumulative Hotfix 15) download URL:

Parcel Repository Location

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1073/parcels/
```

What's new in Cloudera Runtime 7.1.9 SP1 CHF 15

Understand the functionalities and improvements to features of components in Cloudera Runtime 7.1.9 SP1 CHF 15.

Apache Knox

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](#).

Hue

Enhanced session security for Hue

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

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

Apache Ozone

New metrics added in the Grafana dashboard

The following metrics are added in the Grafana dashboard:

- Block count sent from Storage Container Manager to DataNode (Includes both overall and DataNode specific block count)
- Block count received from Ozone Manager
- Command timedout from Storage Container Manager to DataNode

Configuration Changes for the FIPS-compliant SASL Changes

Starting from Cloudera Runtime 7.1.9 SP1 CHF15 onwards, Cloudera supports DIGEST-SHA, a new SASL mechanism, to replace DIGEST-MD5. DIGEST-SHA is similar to DIGEST-MD5 except that DIGEST-SHA uses SHA256 and AES instead of MD5 and DES for message digest and encryption, respectively. As a result, DIGEST-SHA uses only FIPS-compliant algorithms. For more information, see [Step 1: Prepare hosts](#).

Yarn

A new security filtering mechanism introduced inside the Application Master

A new security filtering mechanism inside the Application Master prevents unauthorized mapper or reducer implementations from running within the cluster. This feature ensures that only approved classes or packages are executed inside containers, protecting the cluster from unwanted or unsafe task implementations that may introduce performance, stability, or security risks. For more information, see [Hadoop MapReduce Application Manager based security filter](#)

Fixed issues in 7.1.9 SP1 CHF 15

Know more about the list of fixes that are shipped for CDP Private Cloud Base version 7.1.9 SP1 CHF 15.



Note: Contact Cloudera Support for questions related to any specific hotfixes. The version and build identifier for this CDH Parcel is 7.1.9-1.cdh7.1.9.p1073.76451206.

CDPD-69213: Export/Import, Incremental Export: Tags fail to propagate to target child entities when the source entity is deleted

Previously, when creating tables with multiple levels of depth, a tag applied to the first table was propagated along the lineage. However, after dropping the first table and exporting the entire lineage, child tables did not have the propagated tag after import on the target cluster. This occurred because of a task ordering issue in the deferred actions flow: when deferred actions were enabled during import, the add-propagation task was created before the parent entity was deleted, but ran after the deletion. This caused the system to find the edge between the parent entity and the classification in a deleted state, preventing tag propagation to child entities.

The fix adds a check in the tag propagation logic to skip the edge-status validation when an import is in progress and deferred actions are enabled. This ensures that tags are now correctly propagated to child entities even when the source entity is deleted.

Apache Jira: [ATLAS-5055](#)

CDPD-96779: Spark ORC table column name does not support special characters

Spark threw an error when creating ORC tables with column names containing special characters, such as \$, which was compatible with Hive but not with Spark. This issue is now fixed. Special characters can now be used with Spark.

Apache Jira: [SPARK-32889](#)

CDPD-87921: Oozie SSH action port is not configurable

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

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

Example

```
<port>11100</port>
```

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

Example,

```
"11100" : "oozie.action.ssh.action.port"
```

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

CDPD-93325: Hive2Main.java can print out passwords

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

CDPD-77768: Oozie log parameters are not substituted correctly

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

CDPD-97100: Zeppelin fails after updating Zookeeper and Hadoop components

Previously, updating Zookeeper and Hadoop components caused Zeppelin to fail. This occurred because the updates introduced conflicting versions of the io.netty:netty-transport-native-epoll dependency. The issue is now resolved.

CDPD-94957: Broken lineage on Atlas when calling .cache() or .persist() method

Previously, lineage reporting issues occurred when using the .cache() or .persist() methods on Spark DataFrames, which caused lineage to not be correctly reported into Atlas. This issue is now fixed. The Atlas lineage is no longer wrong in case of cached Dataframes.

CDPD-96543: Missing timeline items for partition delta updaters

Previously, the Catalog Server Operation section in the DDL profile incorrectly displayed that delays in metadata loading were caused by fetching the latest Metastore event ID.

This issue is now fixed by adding the missing timeline items for loading partition names from Metastore to the catalog timeline. These external calls are now accurately tracked to help you find the cause of metadata loading delays.

Apache Jira: [IMPALA-14062](#)

CDPD-96540: Unnecessary metadata reloads for Metastore ALTER_TABLE events

Previously, the catalog service incorrectly triggered file metadata reloads for certain trivial Metastore ALTER_TABLE events.

This issue is now fixed by normalizing null and empty StorageDescriptor parameters to avoid unnecessary reloads. The logs are also improved, when non-trivial StorageDescriptor changes are detected, to display the actual changes.

Apache Jira: [IMPALA-14646](#)

CDPD-96529: Slow metadata operations for wide tables

Previously, when using Ranger authorization, Impala checked column masking policies using the column list. This required loading the table metadata, which triggered unnecessary loading delays, even when running the `INVALIDATE` or `REFRESH` command.

This issue is now fixed. Impala does not depend on the column list to check policies. This prevents unnecessary metadata loading and improves the performance of `INVALIDATE` and `REFRESH` commands on unloaded tables.

Apache Jira: [IMPALA-14703](#)

CDPD-92930: Impala daemon crashes with analytic functions and collections

Previously, running a query with an analytic function on unnested collection columns caused the Impala daemon to crash.

This issue is now fixed by ensuring that data is added to the sorting process only when it is fully available and required.

Apache Jira: [IMPALA-13272](#)

CDPD-91994: Stale query IDs in catalog logs

Previously, Catalog logs for `getPartialCatalogObject` requests displayed incorrect query IDs.

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

Apache Jira: [IMPALA-14494](#)

CDPD-93504: NullPointerException when hive.acid.key.index is missing in ORC files

Previously, a `NullPointerException` (NPE) error occurred when the `hive.acid.key.index` property was missing from an Optimized Row Columnar (ORC) file.

This issue is now resolved. The fix ensures that if the metadata is missing, the system sets the minimum and maximum keys to null and performs a full file scan instead of failing.

Apache Jira: [HIVE-26147](#)

CDPD-94206: Partition column statistics corruption during concurrent updates

Previously, a concurrency issue existed when multiple processes updated partition column statistics at the same time.

This issue is now resolved. The update process is now synchronized to ensure that statistics remain accurate during simultaneous access on the same partition.

Apache Jira: [HIVE-29316](#)

CDPD-94205: Database inconsistency during partition deletion failures

Previously, when a direct SQL processing failure occurred during operations such as dropping partitions, the partial changes made to the database were not reverted before the system attempted an alternative processing method.

This issue is now resolved. The system now sets a savepoint before starting direct SQL processing, ensuring that if a failure occurs, the transaction rolls back to the savepoint before retrying the operation.

Apache Jira: [HIVE-26976](#)

CDPD-93503: Incorrect row order after query-based minor compaction

Previously, query-based minor compaction used an incorrect sorting order for rows in its inner query. This resulted in duplicated rows when you ran the compaction after multiple MERGE statements.

This issue is now resolved. The fix aligns the query-based minor compaction sorting order with the MapReduce compaction order by using the transaction ID, bucket property, and row ID.

Apache Jira: [HIVE-25258](#)

CDPD-93502: Query-based major compaction failure after multiple merge statements

Previously, query-based major compaction failed with a Wrong sort order error when processing tables with multiple Atomicity, Consistency, Isolation, and Durability (ACID) delta directories, such as those created after multiple MERGE statements.

This issue is now resolved. The validation now correctly verifies the row order based on the original transaction ID, bucket property, and row ID, ensuring compatibility with data written by MapReduce and query-based processes. **Apache Jira:** [HIVE-25257](#)

CDPD-93427: Data loss in insert-only tables during query-based minor compaction

Previously, query-based minor compaction produced empty delta files when a table had both aborted and open transactions.

This issue is now resolved. The fix ensures that the compaction process correctly identifies the range of transactions to include, preventing the creation of empty delta files and subsequent data loss.

Apache Jira: [HIVE-29272](#)

CDPD-92365: Unsupported removal of SerDe properties

Previously, you could not remove specific serialization or deserialization (SerDe) properties from an existing Hive table because the UNSET SERDEPROPERTIES clause was not supported.

This issue is now resolved. Hive now supports the UNSET operation for SerDe properties, allowing you to delete specific properties by running the ALTER TABLE <table_name> UNSET SERDEPROPERTIES command.

Apache Jira: [HIVE-21952](#)

CDPD-93500: Multiple YARN UI navigation issues when accessed through Knox

Previously, several issues prevented users from properly accessing YARN-related pages through Knox. This issue is now fixed. The following UI elements and API functions now resolve correctly when proxied through Knox:

- Tracking URL link for running Flink jobs in the YARN UI and YARN UI v2.
- Query parameters in the YARN REST API, such as filtering applications by state.
- Spark History Server links on finished jobs page.

These improvements enhance the user experience when managing and monitoring YARN applications, Flink jobs, and Spark jobs through the Knox gateway.

CDPD-92941: Knox Gateway logging configuration cannot be customized through Cloudera Manager

The Knox Gateway Logging Configuration (Safety Valve) advanced configuration snippet did not apply properly to Knox Gateway. This issue is now resolved by allowing individual logger configurations to be set and modified through the Cloudera Manager UI.

COMPX-14637: Incorrect permissions on NodeManager local directories cause container failures

The NodeManager creates the configured local directories with 755 permissions on startup if the directories did not exist. However, if these permissions were changed after startup or if an administrator created the directories with incorrect permissions before starting YARN, the NodeManager would not reset the permissions, resulting in container failures. This issue is now resolved.

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

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

CDPD-92510: Inaccurate table update timestamps in Hue

Previously, the Data last updated on timestamp for Iceberg tables in Hue displayed incorrect information when tables were altered using Spark. This occurred because the timestamp was sourced from the `transient_lastDdlTime` table property, which does not update correctly during Spark operations. This issue is now fixed by updating the metadata source to use Iceberg snapshot timestamps.

CDPD-94739: Audit logging dependency on debug mode in Hue

Previously, audit logging in Hue only functioned when debug logging was enabled by setting the `DJANGO_DEBUG_MODE` property to true. This caused audit logging to be disabled by default. This issue is now resolved by ensuring audit events are always captured regardless of the debug mode settings. This change ensures Hue compliance and security requirements by maintaining audit logs in all environments.

CDPD-92857: Export All failures in secured multi-tenant clusters

Previously, the Export All feature in Hue failed in secured multi-tenant clusters that used tenant colocation policies. This resulted in access errors during the export process. This issue is now resolved. The Export All functionality now follows tenant boundaries and cluster security policies. However, other known limitations exist. For more information, see [Export All failure in secured multi-tenant clusters](#)

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

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

CDPD-96618: Race condition when registering or unregistering Exception Handling frames

A race condition in libgcc could occur during concurrent registration and EH frames by the memory manager, leading to instability during multi-threaded code generation. This issue is now resolved by ensuring proper padding so that a null-terminated list of FDEs is always provided, eliminating the race condition.

Apache Jira: [KUDU-3545](#)

CDPD-93290: Open Keys summary is not displaying proper data

Previously, the Open Keys summary response was not rendering properly and was not displaying proper data. This issue is now fixed.

CDPD-92660: Recon UI fails to load Disk Usage view with large number of nested directories and keys

Previously, with a large number of entries, the Recon UI **Disk Usage** page failed to load or took an excessive amount of time to populate the tree. This issue is now fixed now. The **Disk Usage** page no longer loads the root path by default. You also get a warning if you try to submit a root path DU request. The root path DU can only be fetched after an explicit submission.

CDPD-93580: Ozone Manager retains native ACL metadata for incoming keys

Previously, if the Ozone client was not updated to the latest version in the user's environment, it sent unwanted ACL metadata during key or file creation. This issue is now fixed. A new server-side configuration ensures that the Ozone Manager removes all native ACL information from the keys before persisting them.

CDPD-94895/CDPD-93951/CDPD-93469: Do not know how many DataNodes cannot accept new writes in a container

Previously, you did not know how many DataNodes were actually readable from the perspective of container allocation and metadata operations. This issue is now fixed and a new metric named

NonWritableNodes is added which shows the number of DataNodes that cannot accept new writes because they are either not in IN_SERVICE and HEALTHY state, cannot allocate new containers or cannot write to existing containers.

CDPD-93860: SpaceUsageSource implementation class Fixed is missing toString()

Previously, the toString() method was missing in SpaceUsageSource.Fixed implementation. The toString() method in SpaceUsageSource.Fixed implementation is used for logging and debugging purposes to display volume space usage information such as capacity, used and available space. This issue is now fixed.

CDPD-93609: Ozone Usage is going beyond 100 percentage

Previously, the Ozone Datanode was using more disk space than allowed for Ozone. This issue is now fixed.

CDPD-88335: DataNode decommissioning fails when other DataNodes are offline due to invalid affinity node in Ratis replication

Previously, DataNode decommissioning failed if one or more DataNodes holding Ratis replicas were already offline. During decommissioning, the Storage Container Manager (SCM) attempts to replicate Ratis container data to maintain redundancy. However, SCM's placement logic still references an affinity node that was previously part of the network topology but has since been removed due to the DataNode going offline. This resulted in repeated placement errors and prevented the decommissioning process from completing successfully. This issue is now fixed.

Apache Jira: [HDDS-13544](#)

CDPD-93901: Not able to track Ozone Storage Container Manager (SCM) safemode rules and exit

In large clusters, the Ozone SCM safemode might take a lot of time to exit. Previously, the you were not able to track the Ozone SCM Manager safemode rules and exit. This issue is not fixed and a new dashboard named "SCM Safemode" is introduced in Grafana which contains a chart for each safemode rule displaying its target and actual value. It also displays if the SCM is in safemode or not by showing "In Safemode" in red and "Exited safemode" in green respectively.

Apache Jira: [HDDS-14039](#)

CDPD-93793: Storage Container Manager (SCM) does not log safemode exit rules at regular intervals

Previously, SCM logged rule statuses at arbitrary time intervals instead of logging them at regular intervals. This issue is now fixed and SCM logs safemode exit rule status at regular intervals while in safemode. The logging interval is configurable through `hdds.scm.safemode.log.interval`. Default value is 1 minute. Each log entry includes overall safemode state, precheck completion status, and current progress of each exit rule.

Apache Jira: [HDDS-14012](#)

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

- [CVE-2023-33202](#) - Bouncy Castle
- [CVE-2025-48734](#) - commons-beanutils
- [CVE-2024-25710](#) - commons-compress
- [CVE-2024-26308](#) - commons-compress
- [CVE-2025-48976](#) - commons-fileupload
- [CVE-2024-47554](#) - commons-io
- [CVE-2021-29425](#) - commons-io
- [CVE-2024-25638](#) - dnsjava
- [CVE-2023-52428](#) - nimbus-jose-jwt

Behavioral changes in Cloudera Runtime 7.1.9 SP1 CHF 15

Behavioral changes for Cloudera Runtime 7.1.9 SP1 cumulative hotfix (CHF) 15 introduce functional adjustments to Cloudera Runtime components.

Hive**HIVE-26976, HIVE-28578, HIVE-29316, HIVE-29030, CDPD-75871: Performance-related issues for HMS**

Previous behavior:

Several performance-related issues for HMS occurred, including stale entries or inconsistent metadata in the database and exceptions during the update operation for multiple table column statistics or for multiple partition-level column statistics.

New behavior:

To resolve these issues, a backend database exclusive row lock is now introduced. The statistics importer tasks are coordinated so that only one importer task that obtains the lock runs at a time. The other tasks to import or update the statistics for the same table or partition continue to be in the pending queue until the current importer task completes and then releases the lock. By default, a blocked or pending importer task retries the lock 100 times in the 3 secs ~ 33 secs duration until a failure occurs.

You can now configure the number of retries and the interval between retries using the following parameters that you can set on the Cloudera Manager Clusters Hive service Configuration tab:

- The `hive.metastore.s4u.nowait.max.retries` parameter allows you to set the number of retries to acquire the row lock immediately without waiting. The default value is 100.
- The `hive.metastore.s4u.nowait.retry.sleep.interval` parameter allows you to set the interval between the retries to acquire a row lock immediately. The default value is 300ms.

Known issues in 7.1.9 SP1 CHF 15

The following section lists the known issues and limitations, the areas of impact, and workaround for each Runtime component.

Known Issues in Atlas

Learn about the known issues in Atlas the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15**CDPD-93307: Hive import logs are truncated during bulk import**

When running the `import-hive.sh` script to perform bulk Hive metadata imports, the logging process fails to retain the expected history. The `import-hive.log` file rotates prematurely when it reaches 10 MB and retains only one backup file (`import-hive.log.1`), instead of the expected 100 MB file size and 20 backup files.

Supply a custom `log4j` configuration file with the correct structure by using the `ATLAS_OPTS` environment variable before running the import script:

```
<?xml version="1.0" encoding="UTF-8" ?>
<!DOCTYPE log4j:configuration SYSTEM "log4j.dtd">
<log4j:configuration xmlns:log4j="http://jakarta.apache.org/log4j/">
  <appender name="console" class="org.apache.log4j.ConsoleAppender">
    <param name="Target" value="System.out"/>
    <layout class="org.apache.log4j.PatternLayout">
      <param name="ConversionPattern" value="%d %-5p - [%t: %x] ~ %m (%C{1}:%L)%n"/>
    </layout>
  </appender>
  <appender name="FILE" class="org.apache.log4j.RollingFileAppender">
```



```

    <param name="File" value="${atlas.log.dir}/${atlas.log.f
ile}"/>
    <param name="Append" value="true"/>
    <param name="maxFileSize" value="100MB"/>
    <param name="maxBackupIndex" value="20"/>
    <layout class="org.apache.log4j.PatternLayout">
      <param name="ConversionPattern" value="%d %-5p - [%t:
%x] ~ %m (%C{1}:%L)%n"/>
    </layout>
  </appender>
  <logger name="org.apache.atlas" additivity="false">
    <level value="info"/>
    <appender-ref ref="FILE"/>
  </logger>
  <logger name="org.apache.kafka.common.config.AbstractConfig"
additivity="false">
    <level value="error"/>
    <appender-ref ref="FILE"/>
  </logger>
</root>
<priority value="warn"/>
<appender-ref ref="FILE"/>
</root>
</log4j:configuration>

```

Export the ATLAS_OPTS variable to point to your custom file before running the import script:

```

export ATLAS_OPTS="-Dlog4j.configuration=file:/path/to/custom-hi
ve-import-log4j.xml"
./import-hive.sh

```

CDPD-87670: [Atlas UI] Apache Atlas Glossary becomes unresponsive when the page size is set to 50

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

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-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-6565: Issue with ddlQueries in Atlas for Hive/Impala tables created in default/non-default db using DAS/HUE respectively

ddlQueries are not created for Impala origin table. ddlQueries created for Impala/Hive CTAS tables have difference in names as compared to same operation for Hive table. The Impala table has the query in the ddlQueries entity name.

None

CDPD-55301: ddlQueries and ALTERNATE_* lineage missing for Spark tables created through spark3-shell

The ddlQueries and ALTERNATE_* lineage missing for Spark tables created using spark3-shell.

None

CDPD-56085: 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

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*.

None

CDPD-58554: Discard audits of specific classification, label, and business metadata

Support to control audits for specific classification, label, and business metadata is not present in the Custom Audit Filters feature.

None

CDPD-40346: The ddlQueries and ALTERNATIVE_ADDCOLS lineage missing for Impala tables.

The ALTERNATIVE_ADDCOLS lineage has some issue when an Impala table is altered and the corresponding lineage is not created.

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

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

None

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

Adding Japanese labels or user defined properties results in the error message: "Invalid label: ###, label should contain alphanumeric characters, _ or -"

None

CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties

Entities, such as hdfs_path, are not returned by the search when their indexable properties are searched using partial search terms made of Japanese or Chinese characters. Only exact matches return results when searching the indexed properties made of Japanese or Chinese characters.

None

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

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

None.

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

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

None.

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

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

None.

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

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

None.

CDPD-43772: Performance issues with Atlas service

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

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

OPSAPS-67783: During rolling upgrade one among two Atlas server failed to start but Cloudera Manager considered as success

Cloudera Manager marks the Execute command Start on service Atlas-1 as a success even when Atlas service had failed to start successfully. In such cases, Atlas logs give the exact reason for Atlas start-up failure.

None

CDPD-19358: "IsIndexable" and "isOptional" value of a typedef's attribute is modified post migration.

In HDP-265, falcon_feed_creation has "stored-in" attribute which has IsIndexable value set to True.

Post migration, "stored-in" is moved to relationshipAttributeDefs and has IsIndexable value set to False.

Similarly, in HDP-265 the table attribute of hive_storage_desc isOptional is set to true. Post migration, isOptional set to False.

None.

CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command

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

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

In converting content from Navigator to Atlas, the create time for Hive tables is not moved to Atlas.

None

CDPD-10574: Suggestion order doesn't match search weights

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

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) does not 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 the metadata collection from HMS. For example, if an ALTER TABLE operation in Spark that is changing a table name and is reporting 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 "@" does not 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, and use 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 does 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 logout 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 CDP.

None.

CDPD-1664: Guest users are redirected incorrectly

Authenticated users logging in to Atlas are redirected to the CDP 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: The 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, to add metadata tags to relationships that you wanted to make sure were unique in the system, you can design the relationship attribute with the property IsUnique equal true. However, in this release, the IsUnique attribute is not enforced.

None.

DOCS-13759: Tag Propagation stops after a certain depth while the lineage is being extended

When a tag is added to an entity at timestamp T1, the entities along the lineage to which the tag must be propagated is calculated at T1. Before tag propagation completes, if the lineage is extended, tag does not propagate to the entities in the extended lineage.

CDPD-41142: When a Kafka console consumer group is run, more than one update audits are seen

After running the console consumer with a consumer group, verify the consumer group entity created, along with the metrics and notifications for the consumer group and topic. The expected result can be: one ENTITY_CREATE audit and one ENTITY_UPDATE audit. But more than one ENTITY_UPDATE audits are seen.

CDPD-40165: Two audits are created for SPARK CTAS table

When following Spark queries are fired:

```
spark.sql("create table table1(id int)")
spark.sql("create table table2 as select * from table1")
```

HMS sends "ENTITY_CREATE" and "ENTITY_FULL_UPDATE_V2".

The extra ENTITY_FULL_UPDATE_V2 message received from HMS is sent as a part of ALTER_TABLE_ADDCOLS event from the HMS Hook side. This behaviour is observed only when the queries are run from Spark SQL and not when the same queries are run from Beeline.

CDPD-39197: Debug metrics returns empty data

When debug metrics is enabled and some operations are performed, the response is empty

CDPD-36495: Updating legacyAttribute from False to True resets the initially created relationshipAttributes values

Creating types, entities, and to start, you must set the relationship with is_legacy_attribute value as **False**.

Later, update the value relationshipDef is_legacy_attribute to True.

For the entities that were created before updating the is_legacy_attribute to True, the relationship Attributes value is reset.

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-29409: Hive import: Suggestion suggests entity which is deleted.

Suggestions suggests tables of a database, which is a deleted entity.

CDPD-25152: Tag propagation through deferred actions consumes additional time as compared to default flow

The additional time might be due to the small overhead added to create / update task vertex and which is run in the background. This also depends on number of tasks queued to be executed.

CDPD-42954: Zeppelin notebook fails after enabling Atlas-HDFS hook

The Zeppelin notebooks are failing with errors after enabling Atlas-HDFS hook in the CDP cluster.

When the below properties are set for atlas-client.properties in Cloudera Manager:

- atlas.jaas.KafkaClient.option.keyTab
- atlas.jaas.KafkaClient.option.principal

Along with adding the properties in /etc/atlas/conf/atlas-application.properties, Cloudera Manager also adds these properties to atlas-application.properties for other services (like Spark).

Adding these properties interferes with the normal flow of the services (like Spark)

To enable HDFS lineage feature, instead of setting these properties through Cloudera Manager, users can manually add the properties directly in /etc/atlas/conf/atlas-application.properties

CDPD-10576: Deleted Business Metadata attributes appear in Search Suggestions

Atlas search suggestions continue to show Business Metadata attributes even if the attributes have been deleted.

None.

CDPD-39427: [HDFS Lineage]: When the input is a directory in case of put/copyfromLocal/cp/mv commands, lineage is not created even though the script succeeds.

When Source is a directory and target is a directory which is already present in Atlas, the command succeeds and inserts the data in the desired location, but lineage is not created.

DOCS-13760: System Attributes search, __classificationNames: Search with parent tag does not return entities associated to its children tags

System attribute search with __classificationNames = parent_tag returns entities associated to parent_tag only and not entities associated to its children tag.

Workaround: Instead of using system attribute, employ the basic search attribute "classification" which lists entities associated with inherited classifications.

CDPD-35818: Basic search with tag filter provides approximateCount as -1 when there is no match and is 0 otherwise

When the following search operations are performed:

- Faceted search with both tag filter and entity filter

The observed approximateCount is -1.

Here both entity filter and tag filter are present and when there is no match the response received is -1.

- Faceted search with only entity filter

Performing a basic search provides an approximate value of 0 when there is no match.

- Faceted search with only tag filter

Whenever there is tag filter in the query and there is no entity match, the approximateCount is -1 and if the tag filter is not available, the response approximateCount is 0

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet `atlas.authentication.method.file` is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export

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

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](#)

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

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

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

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

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

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

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

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

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-88230: Bulk import operations currently do not construct lineage as expected

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

Bulk import operations do not create audits for the imported entities.

Known Issues in Cruise Control

Learn about the known issues in Cruise Control, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15**Rebalancing with Cruise Control does not work due to the metric reporter failing to report the CPU usage metric**

On the Kafka broker, the Cruise control metric reporter plugin may fail to report the CPU usage metric.

If the CPU usage metric is not reported, the numValidWindows in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.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 the [Moving a Host Between Clusters](#).



Note: Cluster nodes affected by this issue are not displayed as unhealthy.

Known Issues in HDFS

Learn about the known issues in HDFS, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15

There are no known issues identified in the releases before Cloudera Runtime 7.1.9 SP1 CHF 14.

Known Issues in Hive

Learn about the known issues in Hive, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15

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

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

The following services support DDL statements:

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

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

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

CDPD-60770: Beeline Authentication Issue with Special Characters in Passwords

When LDAP is enabled, users cannot authenticate with Beeline if the password contains a special character. For example, the following string fails:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>;user=user@XXX;password='R3G#xpXyoy1M0Jb1'
```

Use the -p parameter to execute the Beeline command:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>; -n user@XXX -p 'R3G#xpXyoy1M0Jb1'
```

CDPD-43769: Inconsistent behavior with CHAR comparisons and string literals

In Hive, 'char' data to 'string' data comparison is considered as equal, but 'char' data to 'varchar' data comparison is considered as not equal if the data has trailing spaces. For example,

```
create table test(colchar char(10), colvarchar varchar(20), colstring string);
insert into test values ('a ', 'a ', 'a ');

select count(*) from test where colchar = colstring;
Output returns 1 because char and string data is considered as equal

select count(*) from test where colchar = colvarchar;
Output returns 0 because char and varchar data is considered as not equal
```

It is recommended that you use 'varchar or string' data type instead of 'char' data type for any character or string representation.

CDPD-42726: Wrong results for agg queries from stats after running multi-insert query on managed table

Multiple inserts at the same time into the same partition results in invalid stats. For example:

```
create table source(p int, key int, value string);
insert into source(p, key, value) values (101,42,'string42');

create table stats_part(key int,value string) partitioned by (p int);
from source
insert into stats_part select key, value, p
insert into stats_part select key, value, p;

select count(*) from stats_part;
-- In this scenario, StatsOptimizer helps serving this query because the result should be rowNum of the partition p=101.
-- The result is 1, however, the result should be 2.
```

Rewrite the multi-insert statement by merging the branches inserting into the same partition to one branch.

CDPD-57574: Query execution fails due to NullPointerException in DagUtils.setupQuickStart

Hive queries may fail at runtime with a NullPointerException while executing the Tez graph at DagUtils.setupQuickStart method. A part of the stack trace from the error is shown below:

```
ERROR : Failed to execute tez graph.
java.lang.NullPointerException: null at org.apache.hadoop.hive.ql.exec.tez.DagUtils.setupQuickStart(DagUtils.java:1724)
```

The error usually occurs from an invalid plan that has cycles created by semi-join or map-join edges.

Disable some of the optimizations in order to remove the cycles from the plan. It is sufficient to set one of the following properties:

- set hive.tez.dynamic.partition.pruning=false;
- set hive.tez.dynamic.semijoin.reduction=false;
- set hive.auto.convert.join=false;

To avoid any undesired performance regressions, disable the optimization(s) only for specific queries and not globally (through Cloudera Manager) for the whole workload or cluster.

CDPD-45370: Cartesian join fails with an empty table

When the configuration `hive.tez.cartesian-product.enabled` is set to `true`, query fails if one of the tables involved in a Cartesian join has zero rows. This results in an `IllegalArgumentException` with the message `Number of items is 0. Should be positive`.

To address this issue, set `hive.tez.cartesian-product.enabled` to `false` before running the query.

Apache Jira: [\[HIVE-26582\]](#)

CDPD-43107: SemanticException for INSERT INTO statement when Hive Cost-based Optimizer (CBO) is disabled

If you are running the `INSERT INTO` query and `hive.cbo.enable` is set to `"false"`, the query fails with a `SemanticException`. For example,

```
set hive.cbo.enable=false;

CREATE TABLE mytable (
  id INT,
  str STRING
);

INSERT INTO mytable (id, str) VALUES (1, 'a');
```

```
Output:
org.apache.hadoop.hive.ql.parse.SemanticException: 0:0 Expected
2 columns for insclause-0/default@mytable; select produces 1 col
umns. Error encountered near token ''a''
```

This issue occurs only when columns are specified for the target table in the query. `INSERT INTO mytable values (1, 'a');` does not result in an exception.

Enable CBO and run the query.

CDPD-43957: HiveServer shuts down during replication due to high resource usage

During Hive replication (both bootstrap and incremental), you may notice that the `HiveServer` (HS2) shuts down periodically with the following error:

```
java.sql.SQLException: org.apache.hive.jdbc.ZooKeeperHiveClientE
xception: Unable to read HiveServer2 configs from ZooKeeper
  at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.ja
va:265)
  at org.apache.hive.jdbc.HiveDriver.connect(HiveDriver.java:107)
  at java.sql.DriverManager.getConnection(DriverManager.java:664)
  at java.sql.DriverManager.getConnection(DriverManager.java:247)
  at com.cloudera.enterprise.hive3qt.Hive3QueryTool$HiveOperation.
execute(Hive3QueryTool.java:682)
  at com.cloudera.enterprise.hive3qt.Hive3QueryTool.main(Hive3Qu
eryTool.java:935)
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException:
Unable to read HiveServer2 configs from ZooKeeper
  at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureCo
nnParams(ZooKeeperHiveClientHelper.java:177)
  at org.apache.hive.jdbc.Utills.configureConnParamsFromZooKeeper
(Utills.java:580)
  at org.apache.hive.jdbc.Utills.parseURL(Utills.java:391)
  at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.j
ava:263)
  ... 5 more
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: T
ried all existing HiveServer2 uris from ZooKeeper.
  at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.getServerHosts
(ZooKeeperHiveClientHelper.java:132)
```

```
at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureCon
nParams(ZooKeeperHiveClientHelper.java:172)
... 8 mor
```

This issue occurs when you replicate at scale with an unbalanced cluster setup that has all the roles running on the Cloudera Manager host. As a result, Cloudera Manager ends up in a bottleneck situation because HS2 crashes, preventing further replication. The logs (/var/log/messages) indicate that there were a large number of Java processes running on the host, which exhausted the host's memory and triggered an Out Of Memory Killer process to stop HS2.

You can restart HS2 and reinitiate replication. However, the replication process may take a longer time to complete.

- Follow these recommendations when setting up your source and target clusters:
 - Ensure that no HS2, Hive metastore (HMS), DataNode, or NameNode roles are running on the Cloudera Manager host.
 - Ensure that you have multiple instances of HS2 and HMS roles on different nodes.
 - It is recommended that you have the HS2 and HMS roles on different nodes than the DataNodes or NameNodes.

These practices increase the possibility of Hive replication at scale completing successfully.

- Limit the scale of Hive replication by temporarily disabling replication policies.

CDPD-40730: Parquet change can cause incompatibility

Parquet files written by the parquet-mr library in CDP 7.1.9, where the schema contains a timestamp with no UTC conversion will not be compatible with older versions of Parquet readers. The effect is that the older versions will still consider these timestamps as they would require UTC conversions and will thus end up with a wrong result. You can encounter this problem only when you write Parquet-based tables using Hive, and tables have the non-default configuration `hive.parquet.write.int64.timestamp=true`.

None.

CDPD-41274: HWC + Oozie issue: Could not open client transport with JDBC Uri

Currently only Spark cluster mode is supported in the Oozie Spark Action with Hive Warehouse Connector (HWC).

Use Spark action in cluster mode.

```
<spark xmlns="uri:oozie:spark-action:1.0">
  ...
  <mode>cluster</mode>
  ...
</spark>
```

CDPD-23506: OutOfMemoryError in LLAP

Long running spark-shell applications can leave sessions in interactive Hiveserver2 until the Spark application finishes (user exists from spark-shell), causing memory pressure in case of a high number of queries in the same shell (1000+).

You must close spark-shell so that sessions are closed. Add the owner of the database or the tables as a user with read or read/write access to the tables directly.

CDPD-23041: DROP TABLE on a table having an index does not work

If you migrate a Hive table to CDP having an index, DROP TABLE does not drop the table. Hive no longer supports indexes ([HIVE-18448](#)). A foreign key constraint on the indexed table prevents dropping the table. Attempting to drop such a table results in the following error:

```
java.sql.BatchUpdateException: Cannot delete or update a parent
row: a foreign key constraint fails ("hive"."IDXS", CONSTRAINT "
```

```
IDXS_FK1" FOREIGN KEY ("ORIG_TBL_ID") REFERENCES "TBLS ("TBL_ID"
))
```

There are two workarounds:

- Drop the foreign key "IDXS_FK1" on the "IDXS" table within the metastore. You can also manually drop indexes, but do not cascade any drops because the IDXS table includes references to "TBLS".
- Launch an older version of Hive, such as Hive 2.3 that includes IDXS in the DDL, and then drop the indexes as described in [Language Manual Indexing](#).

Apache Issue: [Hive-24815](#)

CDPD-13636: Hive job fails with OutOfMemory exception in the Azure DE cluster

Set the parameter `hive.optimize.sort.dynamic.partition.threshold=0`. Add this parameter in Cloudera Manager (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)

CDPD-10848: HiveServer Web UI displays incorrect data

If you enabled auto-TLS for TLS encryption, the HiveServer2 Web UI does not display the correct data in the following tables: Active Sessions, Open Queries, Last Max n Closed Queries

Known Issues in Hue

Learn about the known issues in Hue, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15

CDPD-94739: Hue audit logging requires debug mode

In Hue, audit logging is only active when debug logging is enabled by setting `DJANGO_DEBUG_MODE = true`. By default, audit logging is disabled.

Update to the upcoming release to restore audit logging functionality.

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

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark queries in a batch mode launches a shell script Oozie job from Hue and this is not supported on a secure cluster.

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

OPSAPS-69659: Hue service fails on restart with "Unable to find pycpg2 2.5.4" error

Hue service fails to restart and you see the following error: Unable to find pycpg2 2.5.4. This could be because you have installed Python in a non-default location and Hue is unable to locate the pycpg2 PostgreSQL database adapter.

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

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

Key: `PYTHONPATH`

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

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

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

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

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

None.

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

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

None.

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

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

None.

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

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

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

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

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

CDPD-43293: Unable to import Impala table using Importer

Creating Impala tables using the Hue Importer may fail.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

```
[search]
# URL of the Solr Server
solr_url=http://[***HOSTNAME***]:[***PORT***]/solr/
```

For example:

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

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

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

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

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

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

Invalid S3 URI error while accessing S3 bucket

The Hue Load Balancer merges the double slashes (//) in the S3 URI into a single slash (/) so that the URI prefix "/filebrowser/view=S3A:/" is changed to "/filebrowser/view=S3A:". This results in an error when you try to access the S3 buckets from the Hue File Browser through the port 8889.

The Hue web UI displays the following error: “Unknown error occurred”.

The Hue server logs record the “ValueError: Invalid S3 URI: S3A” error.

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.

You should be able to load the S3 browser from both 8888 and 8889 ports.

Alternatively, you can use the Hue server port 8888 instead of the load balancer port 8889 to resolve this issue.

Error while rerunning Oozie workflow

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

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

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

CDPD-17956, CDPD-14167: Setting idle session timeout for Hue does not work when the cluster is secured using Knox SSO

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

None

Query history displays "Result Expired" message

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

None.

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

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

Hue limitation after upgrading from CDH to CDP Private Cloud Base

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

Hue UI is blank upon login after upgrading to CDP 7.1.7 from CDH 6

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

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

1. Log in to Cloudera Manager as an Administrator.
2. Obtain the host name of the Knox Gateway by going to Clusters Knox service Instances .

3. Go to Clusters Hue service Configuration and search for the Knox Proxy Hosts field.



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

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

Unsupported features

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

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

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

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

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

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

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

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

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

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



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

Known Issues in Iceberg

Learn about the known issues in Iceberg, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15

There are no known issues identified in earlier releases.

Known Issues in Impala

Learn about the known issues in Impala, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15

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

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

The following services support DDL statements:

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

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

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

CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the `--tls_ciphersuites` startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.

Ensure that the list in the `--tls_ciphersuites` startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

Apache Jira: [IMPALA-14625](#)

CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work

If you try to run `impala-shell` on RHEL 9 by setting the default python executable available in `PATH` to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, `IMPALA_PYTHON_EXECUTABLE=python3`.

CDPD-81980: Impala query may fail due to backpressure

Impala queries could fail with a dropped due to backpressure error, even when service queue and memory usage appear normal.

Increase the value of the `--mem_limit` startup option.

CDPD-42958: After upgrading the CDH 7.1.9 from CDH 6.x, under certain conditions you cannot insert data into a table

Under the following conditions, after upgrading from CDH 6.x to CDH 7.1.9 you cannot insert data into a table from Impala:

- On CDH 6.x, you created a database with Impala in a user specified HDFS location.
- Using Hive, you then created a table in the database.

Under these conditions, the database and table are stored in the user-specified HDFS directory. After upgrading, the HDFS directory of the table is read-only for Impala. Consequently, from Impala you cannot insert new data into the table because Impala does not have write permission on the HDFS directory.

Workaround: To resolve this issue, use either one of the following workarounds:

- Using the Ranger Web UI, in the policy repository `cm_hdfs`, grant the user 'impala' write permission on the directory where the table resides.
- Enter the following command to grant write permission to user 'impala' on the HDFS directory where the table resides.

```
hdfs dfs -setfacl -m default:user:impala:rwX <HDFS directory>
```

Impala cannot update table if the 'external.table.purge' property is not set to true

Impala cannot update a table using DDL statements if the 'external.table.purge' property is FALSE. ALTER TABLE statements return success with no changes to the table.

ALTER TABLE statements should be issued twice if "external.table.purge" was FALSE initially.



Note: For Iceberg tables, Impala users should always use CREATE TABLE since the Impala CREATE TABLE statement sets the flag to TRUE. However, Impala CREATE EXTERNAL TABLE sets the flag to FALSE.

Impala's known limitation when querying compacted tables

When the compaction process deletes the files for a table from the underlying HDFS location, the Impala service does not detect the changes as the compactions does not allocate new write ids. When the same table is queried from Impala it throws a 'File does not exist' exception that looks something like this:

```
Query Status: Disk I/O error on <node>:22000: Failed to open HDFS file hdfs://nameservice1/warehouse/tablespace/managed/hive/<database>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException: File does not exist: /warehouse/tablespace/managed/hive/<database>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

You must use a single Knox node to access Impala UI.

Impala api calls via kinox require configuration if the kinox customized kerberos principal name is a default service user name

To access impala api calls via kinox, if the kinox customized kerberos principal name is a default service user name, then configure "authorized_proxy_user_config" by clicking Clusters->impala->configuration. Include the kinox customized kerberos principal name in the comma separated list of values `<knox_custom_kerberos_principal_name>=*` where `<knox_custom_kerberos_principal_name>` is the value of the Kerberos Principal in the Knox service. Select Clusters>Knox>Configuration and search for Kerberos Principal to display this value.

CDPD-21828: Multiple permission assignment through grant is not working

None

Problem configuring masking on tables using Ranger

The following Knowledge Base article describes the behavior when we configure masking on tables using Ranger. This configuration works for Hive, but breaks queries in some scenarios for Impala.

For a workaround, see the following Knowledge Base article: [ERROR: "AnalysisException: No matching function with signature: mask\(FLOAT\)" when Impala jobs fail with the following error with signature: mask\(FLOAT\)](#)

IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXXXX"

IMPALA-532: Impala should tolerate bad locale settings

If the LC_* environment variables specify an unsupported locale, Impala does not start.

Add LC_ALL="C" to the environment settings for both the Impala daemon and the Statestore daemon.

Avro Scanner fails to parse some schemas

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

IMPALA-691: Process mem limit does not account for the JVM's memory usage

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the impalad daemon.

To monitor overall memory usage, use the top command, or add the memory figures in the Impala web UI /memz tab to JVM memory usage shown on the /metrics tab.

IMPALA-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Remove trailing semicolon from the Avro schema.

IMPALA-1652: Incorrect results with basic predicate on CHAR typed column

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the RPAD() function to blank-pad literals compared with CHAR columns to the expected length.

IMPALA-1821: Casting scenarios with invalid/inconsistent results

Using a CAST() function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

IMPALA-2005: A failed CTAS does not drop the table if the insert fails

If a CREATE TABLE AS SELECT operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

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

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.

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.

Known Issues in Kafka

Learn about the known issues in Kafka, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15

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

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

The following services support DDL statements:

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

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

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

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker) in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Save Changes > Restart SMM.

The offsets.topic.replication.factor property must be less than or equal to the number of live brokers

The offsets.topic.replication.factor broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a GROUP_COORDINATOR_NOT_AVAILABLE error until the cluster size meets this replication factor requirement.

None

Requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true

The first few produce requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true.

Increase the number of retries in the producer configuration setting retries.

Performance degradation when SSL Is enabled

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with ssl.secure.random.implementation = SHA1PRNG. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

OPSAPS-43236: Kafka garbage collection logs are written to the process directory

By default Kafka garbage collection logs are written to the agent process directory. Changing the default path for these log files is currently unsupported.

None

RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transactional method because we are in an error state
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeFailWithError(TransactionManager.java:1125)
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeAddPartition(TransactionManager.java:442)
    at org.apache.kafka.clients.producer.KafkaProducer.doSend(KafkaProducer.java:1000)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:914)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:800)
    :
    :
    :
    Caused by: org.apache.kafka.common.errors.ClusterAuthorizationException: Cluster authorization failed.
```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting `enable.idempotence` to `false`.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

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

The Debezium Db2 Source connector does not support the evolution (updates) of schemas. In addition, schema change events are not emitted to the schema change topic if there is a change in the schema of a table that is in capture mode. For more information, see [DBZ-4990](#).

None.

CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy compression. The HDFS Stateless Sink connector is only affected if the `Compression Codec` or `Compression Codec` for Parquet properties are set to `SNAPPY`.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: org.apache.hadoop.util.NativeCodeLoader.buildSupportsSnappy()
```

```
Failed to write to HDFS due to java.lang.RuntimeException: native snappy library not available: this version of libhadoop was built without snappy support.
```

Download and deploy missing libraries.



Important: Ensure that you complete steps 1-11 on all Kafka Connect hosts. Additionally, ensure that the advanced configuration snippet in step 12 is configured for all Kafka Connect role instances.

1. Create the `/opt/nativelibs` directory.

```
mkdir /opt/nativelibs
```

2. Change the owner to kafka.

```
chown kafka:kafka /opt/nativelibs
```

3. Locate the directory containing the Hadoop native libraries and copy its contents to the directory you created.

```
cp /opt/cloudera/parcels/CDH/lib/hadoop/lib/native/* /opt/nativelibs
```

4. Verify that `libsnapy.so` was copied to the directory you created.
5. Remove the following from `/opt/nativelibs`.

```
libhadoop.a
```

```
libhadoop.so
libhadoop.so.1.0.0
```

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

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to `requested` instead of `required` and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Unsupported Features

The following Kafka features are not supported in Cloudera Data Platform:

- Only Java and .Net based clients are supported. Clients developed with C, C++, Python, and other languages are currently not supported.
- The Kafka default authorizer is not supported. This includes setting ACLs and all related APIs, broker functionality, and command-line tools.
- SASL/SCRAM is only supported for delegation token based authentication. It is not supported as a standalone authentication mechanism.

- Kafka KRaft in this release of Cloudera Runtime is in technical preview and does not support the following:
 - Deployments with multiple log directories. This includes deployments that use JBOD for storage.
 - Delegation token based authentication.
 - Migrating an already running Kafka service from ZooKeeper to KRaft.
 - Atlas Integration.

Limitations

Collection of Partition Level Metrics May Cause Cloudera Manager's Performance to Degrade

If the Kafka service operates with a large number of partitions, collection of partition level metrics may cause Cloudera Manager's performance to degrade.

If you are observing performance degradation and your cluster is operating with a high number of partitions, you can choose to disable the collection of partition level metrics.



Important: If you are using SMM to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both SMM and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, SMM will not be able to display information about partitions. In addition, Cruise Control will not operate properly.

Complete the following steps to turn off the collection of partition level metrics:

1. Obtain the Kafka service name:
 - a. In Cloudera Manager, Select the Kafka service.
 - b. Select any available chart, and select Open in Chart Builder from the configuration icon drop-down.
 - c. Find \$SERVICENAME= near the top of the display.
The Kafka service name is the value of \$SERVICENAME.
2. Turn off the collection of partition level metrics:
 - a. Go to Hosts Hosts Configuration .
 - b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

Enter the following to turn off the collection of partition level metrics:

```
[KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_entity_update_enabled=false
```

Replace [KAFKA_SERVICE_NAME] with the service name of Kafka obtained in step 1. The service name should always be in lower case.

- c. Click Save Changes.

Known Issues in Knox

Learn about the known issues in Knox that impact or change functionality, and the workarounds in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15**CDPD-76435: Default HTTP client timeout for Streams Messaging Manager requests may cause connection timeouts**

Some Streams Messaging Manager requests, particularly those retrieving topic data, may take up to 30 seconds to complete, even when there are no connectivity issues. Since the default gateway timeout is set to 20 seconds, this can result in connection timeouts despite the system functioning as expected.

Increase the HTTP client and socket timeouts for Streams Messaging Manager services to 60 seconds. This prevents connection timeouts when the requests take longer than the default gateway timeout.

Follow these steps to configure the required properties for the SMM API and SMM UI services:

1. In Cloudera Manager, select the Knox service.
2. Go to Configuration.
3. Find the Knox Simplified Topology Management - cdp-proxy property, and click **+** to add the following lines:

```
SMM-UI:replayBufferSize=512
SMM-UI:httpclient.connectionTimeout=1m
SMM-UI:client.socketTimeout=1m
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

4. Find the Knox Simplified Topology Management - cdp-proxy-api property, and click **+** to add the following lines:

```
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

5. Find the Knox Simplified Topology Management - cdp-proxy-token property, and click **+** to add the following lines:

```
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

6. Save your changes.
7. Refresh the cluster.

Known Issues in Kudu

Learn about the known issues in Kudu, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15**CDPD-94474: Kudu Tablet Server Crash on RHEL9**

Kudu tablet servers kudu-tserver crash on Red Hat Enterprise Linux 9 (RHEL9) x86_64 platforms when the libgcc package is updated to version 11.5.0-11 or newer. This crash manifests as a Segmentation Fault (SIGSEGV) or an Abort (SIGABRT) error under certain workloads.

1. Log in to Cloudera Manager as an Administrator.
2. Navigate to Cloudera Manager Clusters Kudu Configuration .
3. Navigate to the Configuration tab, and add the following value in the Tablet Server Advanced Configuration Snippet (Safety Valve) for gflagfile for the Tablet Server Default Group field:

```
--mrs_use_codegen=false
```

4. Click Save Changes.
5. Restart the Kudu Service.

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.

Workaround: Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

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.

Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

Unsupported features

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 in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF15

No known issues identified in Livy.

Known Issues in Oozie

Learn about the known issues in Oozie, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15

There are no known issues identified in the releases before Cloudera Runtime 7.1.9 SP1 CHF 14.

Known Issues in Ozone

Learn about the known issues in Ozone, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

CDPD-96710: Non-recursive directory deletion fails with the **Directory is not empty** error

S3 clients might display **Directory is not empty** error unexpectedly when recursively deleting a directory through S3 Gateway on Ozone. This is a timing issue. This occurs when the Ozone Manager's internal double buffer has not flushed the child delete transactions to the DB yet.

It typically takes a few seconds for the double buffer to flush under load. So, wait for some time, and then try the recursive delete operation again.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15

CDPD-88335: Datanode decommission fails when other datanodes are offline due to invalid affinity node in Ratis replication

Datanode decommissioning fails if one or more datanodes holding Ratis replicas are offline. During decommission, the Storage Container Manager (SCM) tries to replicate Ratis container data to maintain redundancy. However, SCM's placement logic still references an affinity node that was previously part of the network topology but has since been removed due to the datanode going offline. This causes repeated placement errors and prevents the decommission process from completing successfully.

Apache Issue: [HDDS-13544](#)

CDPD-87429: Unexpected log lines in Ozone snapshot diff output

``ozone.client.server-defaults.validity.period.ms`` configuration introduced to get server default configurations causes Replication Manager to fail when replicating data between Ozone buckets with the following error when the config value does not specify a timeunit.

```

...
java.lang.IllegalStateException: Invalid first line in diff: "25
/07/02 05:58:44 INFO Configuration.deprecation: No unit for ozon
e.client.server-defaults.validity.period.ms(3600000) assuming MI
LLISECONDS"
at com.cloudera.enterprise.distcp.OzoneCopyListingUtils.parseOz
oneDiff(OzoneCopyListingUtils.java:397)
...

```

Overwrite the ``ozone.client.server-defaults.validity.period.ms`` to configure the value to have time unit (ns, ms, s, m, h). For example, 3600000ms.

CDPD-71648: Snapshot creation failing with "PERMISSION_DENIED" error on a bucket created during pre-upgrade in 7.1.7 SP3

On upgrade of cluster from 7.1.7 SP3 to 7.1.9 SP1, we are not able to create snapshots on a pre-upgrade volume or bucket with "PERMISSION_DENIED" error, since in 7.1.7SP3, we were using full KRB owner name with domain but currently, we only use and search for user shortname.

Update the owner of the bucket with the user shortname and then create snapshot by running the commands:

```

ozone sh bucket update --user=[***USER_SHORT_NAME***] {volume}/{
bucket}
ozone sh snapshot create {volume}/{bucket}
[***SNAPSHOT_NAME***]

```

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

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

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

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

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

CDPD-64865: Intermittent timeout in TestBlockDeletion.testBlockDeletion (BlockDeletionService stuck)

If the event of SCM exiting safe mode was triggered multiple times, this might eventually cause the SCMBlockDeletingService to transition into the PAUSING state.

Check the SCM logs and ensure that the SCMBlockDeletingService is not running, and then restart the SCM.

CDPD-81569: Recon is not reflecting correct metrics for Volumes, Buckets, Keys, and Pending Key Deletions

Recon is displaying data for container and space utilisation only. The metrics for Volumes, Buckets, Keys, and Pending Key Deletions are not being updated and displayed correctly.

Apache Jira: [HDDS-12607](#)

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: 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.

The Prometheus binary is not available in CDP Private Cloud Base 7.1.9 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.

After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and if Ozone is in the Non-HA environment, an exception message is observed during the finalization of the Ozone upgrade.

During the finalization of the upgrade, a `ClassNotFoundException` message for `org.cloudera.log4j.redactor.RedactorAppender` class may be displayed. The error message can be ignored because the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

CDPD-68951: In CDP Private Cloud Base 7.1.9 CHF2 version and lower, the command `ozone sh key list <bucket_path>` displays the `isFile` flag in a key's metadata as `false` even when the key is a file. This issue is rectified in CDP Private Cloud Base 7.1.9 CHF3. However, the pre-existing (pre-upgrade) key's metadata cannot be changed.

None

CDPD-67771: When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(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.

CDPD-64398: In SCM with non-HA configuration, Secret key manager does not initialise. Hence, the startup of OM and DataNode fails because it cannot get a secret key. This key is used when security is enabled and used in Block token and container token verification while communicating between the Ozone client, OM, and Datanode.

This issue does not occur in SCM with HA configuration.

You must force exit from the safe mode for SCM. This triggers the initialization of the Secret key manager. Below are the options:

- exit safe mode manually whenever SCM is started
- disable safe mode by setting `hdds.scm.safemode.enabled=false` in safety valve for SCM configuration

Impact of disabling safe mode and the purpose of safe mode:

- Write should not fail once the cluster is out of safe mode
- Read of existing data should not fail after the cluster is out of safe mode
- Unnecessary re-replication should be avoided during cluster restart

So disabling safe mode does not have any major impact on Ozone.

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.

CDPD-52412: keyManagerImpl#listStatus exceeds the maximum RPC length.

To reduce the payload, you must increase the part size.

OPSAPS-68159: If you did not deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Error when distributing parcel to host, Ignoring non-compliant parcel manifest error is displayed after Cloudera Manager is upgraded to 7.11.3.

If you encounter the error, perform the following steps:

1. Deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.11.3.
2. Restart the cluster with a delay of 10 minutes.
3. Continue to upgrade CDH 7.1.8 to CDH 7.1.9.

OPSAPS-68159: If you did not deactivate the Ozone parcel 718.2.x on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Ozone roles during the CDH 7.1.8 upgrade to CDH 7.1.9.

If you encounter the error, perform the following steps:

1. Deactivate the Ozone parcel 718.2.x.
2. Restart the Ozone service.
3. Perform Finalize Upgrade for the Ozone service.

Step result: The Ozone roles will be displayed in green.

CDPD-60989: The packaging version for Apache Ozone points to the 1.2.0 older version. This is a version string problem and not a packaging issue. The version of the Apache Ozone binary is closest to 1.4.0.

None. This only affects the JAR versioning.

CDPD-60366: Native library loader fails when system property `native.lib.tmp.dir` is not set. It fails because the library is copied to / instead of the cwd.

Setting `native.lib.tmp.dir` to a certain path like `/tmp` should solve the issue. To set this system property, you must add `-Dnative.lib.tmp.dir=/tmp` to `ozone_java_opts` in Cloudera Manager configuration of the Ozone cluster.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-60598: DataNode can consume up to the JVM heap size worth of direct memory buffers under specific failure scenarios. This can lead to a DataNode crash. This is due to Netty allocating and not freeing the direct memory buffers.

Restarting DataNode will resolve this issue.

CDPD-60466: The Ozone client is missing in the Cloudera Manager's Service Monitor node. This causes Ozone Canary's health check to fail.

Ensure you install Ozone CLI on the same node as Cloudera Manager's Service Monitor.

CDPD-60012: The Ozone SCM may be stuck in safe mode after upgrade to 7.1.9. This is due to a bug in how the SCM accounts for open storage containers.

The workaround is to temporarily set the Ozone configuration `hdds.scm.safemode.threshold.pct` to a lower value like 0.90 and restart the SCM.

CDPD-59679: In some scenarios, excess UNHEALTHY replicas of EC containers may not be removed from the cluster.

None

CDPD-50116: Topology-aware reads can provide better performance by routing applications to the closest replica of a file, if available, on the same physical node or same rack. However this feature is disabled by default in CDP 7.1.9.

This feature can be enabled by setting `ozone.network.topology.aware.read` to true using Cloudera Manager safety valve.

CDPD-57165: The DataNode disk usage thread may abort due to Ratis log tailing.

Use `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory` to calculate disk usage. Configure in the `ozone-site.xml` property: `hdds.datanode.du.factory` value: `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory`

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.

CDPD-55513: Hive external table replication policy does not migrate Hive tables in HDFS storage to Ozone.

To replicate data from HDFS to Ozone, see [Migrating your data from HDFS to Ozone](#). You can use HMS Mirror to replicate the metadata.

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"`

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_SHAREDLIB_DIR=/VAR/TMP`
3. Restart Ozone.

CDPD-49137: Ozone Manager Kerberos token expires for SCM communication and OM does not log in again.

Sometimes, OM's Kerberos token is not updated and it stops to communicate with SCM. When this occurs, writes start failing.

Restart OM or set the safety valve `hadoop.kerberos.keytab.login.autorenewal.enabled = true`

CDPD-56684: Keys get deleted when you do not have permission on volume

When a volume is deleted, it recursively deletes the buckets and keys inside it and only then deletes the volume. The volume delete ACL check is done only in the end, due to which you may end up deleting all the data inside the volume without having delete permission on the volume.



Note: There was no bucket/key permission available which allowed the user to delete them recursively.

CDPD-50610: Large file uploads are slow with OPEN and stream data approach

Hue file browser uses the append operation for large files. This API is not supported by Ozone in 7.1.9 and therefore large file uploads can be slow or timeout from the browser.

Use native Ozone client to upload large files instead of the Hue file browser.

OPSAPS-64097: Ozone service restart failed at SCM

Stopping SCM service using Cloudera Manager can sometimes timeout and need a retry. The Cloudera Manager API waits for 90 seconds which is not sufficient under certain circumstances.

Retry the shutdown using Cloudera Manager if the SCM still shows as running after a refresh of the service status.

OPSAPS-66469: Ozone-site.xml is missing if the host does not contain HDFS roles

The client side `ozone-site.xml` (/etc/hadoop/conf/ozone-site.xml) is not generated by Cloudera Manager if the host does not have any HDFS role. Because of this, issuing Ozone commands from that host fails because it cannot find the service name to host name mapping. The error message is similar to this: `# ozone sh volume list o3://ozoneabc 23/03/06 18:46:15 WARN ha.OMProxyInfo: OzoneManager address ozoneabc:9862 for serviceID null remains unresolved for node ID null` Check your `ozone-site.xml` file to ensure ozone manager addresses are configured properly.

Add the HDFS gateway role on that host.

OPSAPS-67607: Cloudera Manager FirstRun failure at the “Upload YARN MapReduce Framework JARs” step.

If this failure is attributed to the broken symbolic link, `/var/lib/hadoop-hdfs/ozone-file-system-hadoop3.jar`, it is likely due to the presence of the user `hdfs` on the node prior to CDP parcel activation. As a result, the Cloudera Manager agent skips the initialization related to HDFS, leading to the non-creation of the `/var/lib/hadoop-hdfs` directory.

Create the directory `“/var/lib/hadoop-hdfs”` on all nodes followed by the deactivation and activation of the CDP parcel (deactivate and activate the Ozone parcel instead, in case Ozone parcel is used).

CDPD-50447: When SCM High Availability is enabled, each of the SCM web UIs report the host of the web UI as the leader of HA, and the other two as followers. This gives wrong information

Correct output is available by running the `ozone admin scm roles --service-id=<ID>` command.

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](#).

OPSAPS-62327: In an Ozone cluster without any gateway roles, Ozone is unable to deploy client configurations and displays the ConfigGenException error.

You must add the Ozone gateway roles to the cluster.

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-36389: The configurations `datanodes.involved.max.percentage.per.iteration` and `size.moved.max.per.iteration` are meant to limit the max number of DataNodes that is involved and maximum size that can move in an iteration. This bug causes balancer to stop an iteration when it is two DataNodes or 1 Container size (5GB) away from hitting these limits. However, these DataNodes can again be considered for balancing in the next iteration. This means the cluster is balanced after enough iterations, albeit a bit slowly. This bug is apparent in small clusters of around four DataNodes where the DataNodes can be either the source or target for a lot of moves but the iteration gets stopped when three DataNodes have been involved. It takes a higher number of iterations to eventually balance this cluster. While this is a performance issue, it does not prevent balancer from ultimately balancing the cluster.

Increase the speed for balancing by decreasing the interval between each iteration using the configuration `balancing.iteration.interval`. The value of this configuration must be greater than `hdds.datanode.du.refresh.period`. `size.moved.max.per.iteration` can be increased to allow more data to move in one iteration.

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-35141: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2> (state=08S01,code=40000) java.sql.SQLException: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2>. This issue does not occur if the source and target buckets are different in Hive queries. Currently, copying across the same bucket is only supported.

Avoid different buckets in source and target path. Ensure that `hive.exec.stagingdir` is unset or set to the same bucket as the warehouse so that staging directories are created in the same bucket as the warehouse.

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-41184: With Legacy buckets, FileSystem Optimised is not interoperating with the Ozone shell command. The cause is that the directory key entry in the DB KeyTable is stored as "dir1/" with trailing slash. But while performing the described operation, Ozone shell (o3://) is normalizing the given path and removing the trailing slash "/" from it. That results in KEY_NOT_FOUND exception.

There are three workarounds:

- Use FileSystem API to Delete the Directories rather than Shell-Command API.
- Use FSO buckets instead of Legacy Buckets. As in FSO, you can create Intermediate Directories and Delete Directories using the Ozone shell commands.

- Disable and set the configuration `ozone.om.enable.filesystem.pathsflag` to false in order to delete the directories. This is generally not a preferred workaround because the cluster must be restarted to pick up the new changes.

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-12542: Mount of Ozone filesystem with the help of FUSE fails.

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-42897: EC writes are failing with "No enough datanodes to choose" after EC replication config is set globally.

EC writes start failing when large number of pipelines are created as a result of multiple EC configs with different chunk sizes used to write keys.

If standard EC configs (that is, rs-3-2-1024k) are used to write keys, number of pipelines created per DataNode is limited to five and this issue is not seen with standard EC configs.

The recommendation is not to create too many random chunk sizes. It is configurable because, users can decide based on their workload. But do not to have separate chunk sizes for each file.

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-40560: Filesystem Operations using Hadoop S3a connector on a FILE_SYSTEM_OPTIMIZED bucket fails. org.apache.hadoop.ozone.om.exceptions.OMException: Unable to get file status: volume: s3v bucket: fso key: test/

Do not run Hadoop S3a commands on a FILE_SYSTEM_OPTIMIZED bucket. Use OBJECT_STORE bucket layouts.

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.

CDPD-43432: Ozone Service in fault state in DataNode - Long Running setup.

Upgrade RocksDB to the latest version.

OPSAPS-63999: In the newly installed cluster, the Finish upgrade option is clickable.

None

OPSAPS-64648: Failed to start Ozone node using Cloudera Manager if default log path /var/log/hadoop-ozone does not exist. If this path does not exist, any Ozone nodes (for example SCM or DataNode) restart will fail.

Run the following command `sudo -u hdfs mkdir -p /var/log/hadoop-ozone` or replace hdfs with the user Ozone roles that are running.

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.

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 in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15

OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager

Cloudera Manager does not display a Log Files menu for SMM UI role (and SMM UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by SMM UI is not supported by Cloudera Manager.

View the SMM UI logs on the host.

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in SMM safety valve in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server
Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Add the
following value for bootstrap servers>Save Changes > Restart SMM.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

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

KNOX does not send neither the Authentication header nor the hadoop.auth cookie. Because of this, the SMM UI returns an HTTP 401 response and sets the www-authenticate: Negotiate header. After this, KNOX still does not send the Authentication header. This results in the SMM UI being inaccessible through Knox.

You can access the SMM UI directly in Cloudera Manager at Clusters SMM Streams Messaging Manager Web UI and log in using the Kerberos username and password.

CDPD-82560: Streams Messaging Manager known unprotected endpoint

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

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

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 Spark

Learn about the known issues in Spark the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15

CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the RuntimeWarning: Failed to add file message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

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

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

The following services support DDL statements:

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

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

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

CDPD-67517: Spark3 tests are failing if /tmp is mounted as NOEXEC.

Map the tmpdir to a writable path in spark3-conf/spark-defaults.conf using the following steps:

1. Go to the Cloudera Manager.
2. Go to the Spark service.
3. Click the Configuration tab.
4. Select Scope Gateway .
5. Select Category Advanced .
6. Locate the Spark Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-defaults.conf_client_config_safety_valve property.
7. Map the tmpdir to a writable path:

```
spark.driver.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
spark.executor.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
```

8. Enter a Reason for change, and then click Save Changes to commit the changes.
9. Deploy the client configuration.

CDPD-23817: In the upgraded Cluster, the permission of /tmp/spark is restricted due to the HDP configuration hive.exec.scratchdir=/tmp/spark.

If you are using the /tmp/spark directory in the CDP cluster, you must provide the required additional Policies/ACL permissions.

CDPD-22670 and CDPD-23103: There are two configurations in Spark, "Atlas dependency" and "spark_lineage_enabled", which are conflicted. The issue is when Atlas dependency is turned off but spark_lineage_enabled is turned on.

Run Spark application, Spark will log some error message and cannot continue. That can be restored by correcting the configurations and restarting Spark component with distributing client configurations.

CDPD-23007: Mismatch in the Spark Default DB Location. In HDP 3.1.5, hive_db entities have one attribute - 'location' which is configured to the '/managed' path. In fresh install of CDP 7.1.7, hive_db entities now have 2 attributes 'location' configured to '/external' path and 'managedLocation' configured

to `/managed` path. In the AM2CM migration (HDP 3.1.5 -> CDP 7.1.7), the `'location'` attribute from `hive_db` entities in HDP 3.1.5 comes unaltered to CDP 7.1.7 and hence maps to `/managed` path.

This issue arises only if you are upgrading from HDP 3.1.5 to CDP 7.1.7. If you are performing a fresh install of CDP 7.1.7, you can ignore this issue.

None

CDPD-217: The Apache Spark connector is not supported

The old *Apache Spark - Apache HBase Connector* (shc) is not supported in CDP releases.

Use the new HBase-Spark connector shipped in CDP release.

CDPD-3038: Launching pyspark displays several HiveConf warning messages

When pyspark starts, several Hive configuration warning messages are displayed, similar to the following:

```
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of name hive.meta
store.runworker.in does not exist
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of
name hive.masking.algo does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of
name hive.metastore.runworker.in does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of na
me hive.masking.algo does not exist
```

These errors can be safely ignored.

Known Issues in Sqoop

Learn about the known issues in Sqoop, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15

CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version

If you have upgraded to CDP PvC Base 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to CDP PvC Base 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see [Troubleshooting Apache Sqoop issues](#).

CDPD-44431: Using direct mode causes problems

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the `--direct` option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the `sqoop.enable.deprecated.direct` property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through `-Dsqoop.enable.deprecated.direct=true`.

CDPD-3089: Avro, S3, and HCat do not work together properly

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

Parquet columns inadvertently renamed

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.

PARQUET-99: Importing Parquet files might cause out-of-memory (OOM) errors

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.

Known Issues in Schema Registry

Learn about the known issues in Schema Registry, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15**CDPD-40380: Authorization checking issue when Kerberos is disabled**

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the `ranger.plugin.schema-registry.service.name` field.
3. Replace `GENERATED_RANGER_SERVICE_NAME` with the actual name of the service.
4. Restart the Schema Registry service.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

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

Schema Registry uses a schema fingerprinting mechanism to differentiate between schemas. However, fingerprinting does not take into consideration the schema attributes of the field type. As a result, if you have two schemas where the only difference is that one has type attributes defined and the other does not, they will be considered identical by Schema Registry. For example, the following schemas are considered identical:

```
#Schema V1
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"local_timestamp_micros_long","type":"long"}]}

#Schema V2
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"local_timestamp_micros_long","type":{"type":"long","logicalType":"local-timestamp-micros"}}]}
```

Notice that the only difference is that in the second schema, the `local_timestamp_micros_long` field has a logical type specified. In cases like this, the new version of the schema is not created, the initial version is used. This is true even if the data that is being produced has a new schema version. The ID of the first schema version is used and is put in the serialized record. The new schema version is not created.

This issue is common when using change data capture (CDC) connectors like the Debezium connectors. This is because CDC connectors create schemas with the logical type decimal based on the column type in the database schema. For example:

```
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [{"name": "database_column", "type": {"type": "bytes", "logicalType": "decimal", "precision": 64, "scale": 0}}]}
```

If the database schema changes (for example, the column type), it is possible that only scale changes, which is a schema attribute.

```
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [{"name": "database_column", "type": {"type": "bytes", "logicalType": "decimal", "precision": 64, "scale": 1}}]}
```

In this case, even though scale changed to 1, the first version of the schema is used where scale is 0. As a result, the data is consumed with the wrong scale.

Avoid using logical types or other attributes. Alternatively, ensure that there are no changes in the logical types or other attributes between schema versions.

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

Schema Registry does not start if the address specified in the Load Balancer Address Ranger property does not end with a trailing slash (/).

Set the value of the RANGER_REST_URL Schema Registry environment variable to an address that includes a trailing slash.

1. In Cloudera Manager, select the Schema Registry service.
2. Go to Configuration.
3. Find the Schema Registry Server Environment Advanced Configuration Snippet (Safety Valve) property and add the following:

```
Key: RANGER_REST_URL
Value: [***RANGER REST API URL***]
```

Replace [***RANGER REST API URL***] with an address that can be used by Schema Registry to access Ranger. Ensure that the address ends with a trailing slash. For example: `http s://ranger-1.cloudera.com:6182/`

4. Restart the Schema Registry service.

OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

1. Access the Ranger Console (Ranger Admin web UI).
2. Click the `cm_atlas` resource-based service.
3. Add the `schemaregistry` user to the `all - *` policies.
4. Click `Manage Service Edit Service`.
5. Add the `schemaregistry` user to the `default.policy.users` property.

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

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect

instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Known Issues in Streams Replication Manager

Learn about the known issues in Streams Replication Manager, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15

CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

CDPD-11079: Blacklisted topics appear in the list of replicated topics

If a topic was originally replicated but was later excluded for replication, 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, this topic will be included in the response. Additionally, the excluded topic will also be visible in the SMM UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

CDPD-30275: SRM may automatically re-create deleted topics on target clusters

If `auto.create.topics.enable` is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with `srm-control`. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1]
```

2. Wait until SRM is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

Limitations

SRM cannot replicate Ranger authorization policies to or from Kafka clusters

Due to a limitation in the Kafka-Ranger plugin, SRM cannot replicate Ranger policies to or from clusters that are configured to use Ranger for authorization. If you are using SRM to replicate data to or from a cluster that uses Ranger, disable authorization policy synchronization in SRM. This can be achieved by clearing the Sync Topic Acls Enabled (`sync.topic.acls.enabled`) checkbox.

SRM cannot ensure the exactly-once semantics of transactional source topics

SRM data replication uses at-least-once guarantees, and as a result cannot ensure the exactly-once semantics (EOS) of transactional topics in the backup/target cluster.



Note: Even though EOS is not guaranteed, you can still replicate the data of a transactional source, but you must set `isolation.level` to `read_committed` for SRM's internal consumers. This can be done by adding `[***SOURCE CLUSTER ALIAS***]->[***TARGET CLUSTER ALIAS***].consumer.isolation.level=read_committed` to the Streams Replication Manager's Replication Configs SRM service property in Cloudera Manager. The `isolation.level` property can be set on a global connector or replication level. For example:

```
#Global connector level
connectors.consumer.isolation.level
=read_committed
#Replication level
uswest->useast.consumer.isolation.level
evel=read_committed
```

SRM checkpointing is not supported for transactional source topics

SRM does not correctly translate checkpoints (committed consumer group offsets) for transactional topics. Checkpointing assumes that the offset mapping function is always increasing, but with transactional source topics this is violated. Transactional topics have control messages in them, which take up an offset in the log, but they are never returned on the consumer API. This causes the mappings to decrease, causing issues in the checkpointing feature. As a result of this limitation, failover operations for transactional topics is not possible.

Known Issues in YARN 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 in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15

COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

None.

COMPX-4644: Queue capacity rounding problem when configuration is initially set via YARN

When setting the capacity scheduler configuration through the YARN/Cloudera Manager configuration, there may be capacity values that use multiple decimal places. This results in rounding/floating point precision discrepancies in the UI when trying to validate that all sibling capacities equal 100%. The UI looks like all the numbers add up to 100, but the validation still displays an error and does not allow to save the capacities. It is also observed that the capacity is being calculated as, for example, 99.9999999991 in the backend.

- Create queues within the UI, or
- Ensure that capacities configured through the Capacity Scheduler safety valve do not have more than one decimal place.

20202 Database migration after enabling opt-in migration

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to CDP 7.1.9, forcing migration from H2 to PostgreSQL database.

- Upgrade to CDP 7.1.9 CHF2, moving back to H2 database.
- Upgrade to CDP 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version CDP 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

CDPD-56559: MapReduce jobs can intermittently fail during a rolling upgrade.

During a rolling upgrade between CDP versions 7.1.8 and 7.1.9, MapReduce jobs may fail with message, `RuntimeException: native snappy library not available`. Although the native Snappy compression library is not loaded, a checkmark displays to indicate that the Snappy compression library is loading for NodeManagers that are pending upgrades. This causes the MapReduce jobs that are associated with the NodeManagers to fail. After the upgrade, the jobs work as expected. This issue only impacts rolling upgrades from before CDP 7.1.9 to a higher version.

None.

COMPX-12021 Queue Manager configurations on Scheduler Configuration page are not working

When setting the following properties on the YARN Queue Manager UI, the properties are set in the `capacity-scheduler.xml` file which does not have any effect on YARN. The properties need to be set in the `yarn-site.xml` file, which does not happen when you set them through YARN Queue Manager.

- "Maximum Application Priority" – `"yarn.cluster.max-application-priority"`
- "Enable Monitoring Policies" – `"yarn.resourcemanager.scheduler.monitor.enable"`
- "Monitoring Policies" – `"yarn.resourcemanager.scheduler.monitor.policies"`



Note: You can also set this property on the YARN configuration page in Cloudera Manager as "Capacity Scheduler Preemption".

- "Preemption: Observe Only" – `"yarn.resourcemanager.monitor.capacity.preemption.observe_only"`
- "Preemption: Monitoring Interval (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.monitoring_interval"`
- "Preemption: Maximum Wait Before Kill (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.max_wait_before_kill"`
- "Preemption: Total Resources Per Round" – `"yarn.resourcemanager.monitor.capacity.preemption.total_preemption_per_round"`
- "Preemption: Over Capacity Tolerance" – `"yarn.resourcemanager.monitor.capacity.preemption.max_ignored_over_capacity"`
- "Preemption: Maximum Termination Factor" – `"yarn.resourcemanager.monitor.capacity.preemption.natural_termination_factor"`
- "Enable Intra Queue Preemption" – `"yarn.resourcemanager.monitor.capacity.preemption.intra-queue-preemption.enabled"`

1. In Cloudera Manager, select the YARN service.
2. Click the Configuration tab.
3. Search for `yarn-site.xml`.
4. Under YARN Service Advanced Configuration Snippet (Safety Valve) for `yarn-site.xml`, add the corresponding parameter and value you need.
5. Click Save Changes.
6. Restart the YARN services.

COMPX-13177: QueueManager webapp requests fail with 'HTTP ERROR 400

`java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA'`

Products:

- Cloudera Manager for CDP Private
- Cloud Base Cloudera Manager for CDP Public Cloud

Context:

- Centos 7.8 and Redhat 7.8 operating systems, when FIPS support is enabled.

Problem:

- When attempting to display the Yarn Queue Manager interface, Cloudera Manager displays an error: "HTTP ERROR 400 java.net.ConnectException: Unsupported cipher suite 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.

Third-party applications do not launch if MapReduce framework path is not included in the client configuration

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the `mapreduce.application.framework.path` property is set to the appropriate value, but third-party applications with their own configurations does not launch.

Set the `mapreduce.application.framework.path` property to the appropriate configuration for third-party applications.

JobHistory URL mismatch after server relocation

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated the `mapred-site.xml` file that references the correct JobHistory Server.

YARN cannot start if Kerberos principal name is changed

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN does not start. In such cases, the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the `znode` and restart the YARN service.
- Use the `reset ZK ACLs` command. This also sets the `znodes` below `/rmstore/ZKRMStateRoot` to `world:anyone:cdwra` which is less secure.

Queue Manager does not open on using a custom user with a default Kerberos principal

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager `process_username` property matches the YARN `process_username` property.

Known Issues in Zeppelin

Learn about the known issues in Zeppelin the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF15**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.

BUG-125263: Zeppelin service move fails on clusters upgraded from HDP3.1.5

Resolve the circular symlink issue on the Zeppelin node by linking the conf directory to a new directory under /etc/zeppelin:

- # mkdir -p /etc/zeppelin/<version>/0
- # rm /usr/hdp/<version>/zeppelin/conf
- # ln -s /etc/zeppelin/<version>/0 /usr/hdp/<version>/zeppelin/conf

Where version is the HDP version. For example, 7.1.x-yyy. Restart the Zeppelin server in Ambari.

CDPD-3090: Due to a configuration typo, functionality involving notebook repositories does not work

Due to a missing closing brace, access to the notebook repositories API is blocked by default.

From the CDP Management Console, go to Cloudera Manager for the cluster running Zeppelin. On the Zeppelin configuration page (Zeppelin serviceConfiguration), enter shiro urls in the Search field, and then add the missing closing brace to the notebook-repositories URL, as follows:

```
/api/notebook-repositories/** = authc, roles[{{zeppelin_admin_group}}]
```

Click Save Changes, and restart the Zeppelin service.

CDPD-2406: Logout button does not work

Clicking the Logout button in the Zeppelin UI logs you out, but then immediately logs you back in using SSO.

Close the browser.

OPSAPS-59802: Zeppelin and Livy roles should be co-located on the same host.

When installing or upgrading to CDP Private Cloud Base, you must co-locate all Zeppelin and Livy roles on the same cluster host due to an issue with certificate generation.

Known Issues in ZooKeeper

Learn about the known issues in ZooKeeper the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 15.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 15

There are no new known issues identified in the previous releases.

Cumulative hotfix 14

Know more about the cumulative hotfix 14 for CDP 7.1.9 SP1. This cumulative hotfix was released on January 08, 2026.

Cloudera Runtime 7.1.9.1069-3 (Cumulative Hotfix 14) download URL:**Parcel Repository Location**

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1069/parcels/
```

Fixed issues in 7.1.9 SP1 CHF 14

Know more about the list of fixes that are shipped for CDP Private Cloud Base version 7.1.9 SP1 CHF 14.



Note: Contact Cloudera Support for questions related to any specific hotfixes. The version and build identifier for this CDH Parcel is 7.1.9-1.cdh7.1.9.p1069.74330271.

CDPD-94496: Security vulnerability related to Apache Tika

This fix addresses a critical XML External Entity (XXE) vulnerability, as described in [CVE-2025-66516](#), in the Apache Tika core and parser modules. This vulnerability could be exploited by using a crafted XFA file within a PDF document.

Backport of Apache Tika version

Apache Tika 2.9.4 is now forked, patched and built internally into 2.9.4.cldr-b21. This applies for NiFi and CFM as well.

Cloudera Search uses the downstream version of Tika 2.4.1 (2.4.1.cldr-b12) as well which is essentially patched with the CVE fix.

HIVE-26456: Duplicate rows after query-based minor compaction

Previously, duplicate rows appeared in table results after you ran a minor compaction.

This issue is resolved by updating the sorting process to match standard compaction methods, ensuring data is processed in the correct order.

Apache Jira: [HIVE-26456](#)

HIVE-26147: Null pointer exception when ACID key index is missing

Previously, queries failed with a null pointer error when reading ACID-compliant ORC files that were missing specific index metadata.

This issue is resolved by allowing the system to perform a full file scan when the index metadata is missing instead of failing.

Apache Jira: [HIVE-26147](#)

HIVE-21052: Cleanup failure for aborted transactions

Previously, if a transaction aborted before partition registration, the system flagged it as empty.

This issue is resolved by using a marker in the transaction components to track early-stage writes, ensuring the cleaner identifies and removes aborted data directories.

Apache Jira: [HIVE-21052](#)

HIVE-24374: Metadata removal for aborted writes in batched delta files

This issue is now fixed.

Apache Jira: [HIVE-24374](#)

HIVE-25502: Data loss when cleaning aborted transactions with dynamic partitions

Previously, running the cleaner after an aborted streaming transaction could cause data loss in unrelated partitions.

This issue is now fixed.

Apache Jira: [HIVE-25502](#)

HIVE-25257: Incorrect row order validation for query-based major compaction

Previously, query-based major compaction failed when processing tables with multiple transaction statements, such as consecutive MERGE operations.

This issue is now fixed.

Apache Jira: [HIVE-25257](#)

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

- [CVE-2025-66516](#) - Apache Tika

Known issues in 7.1.9 SP1 CHF 14

The following section lists the known issues and limitations, the areas of impact, and workaround for each Runtime component.

Known Issues in Atlas

Learn about the known issues in Atlas the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

CDPD-93307: Hive import logs are truncated during bulk import

When running the import-hive.sh script to perform bulk Hive metadata imports, the logging process fails to retain the expected history. The import-hive.log file rotates prematurely when it reaches 10 MB and retains only one backup file (import-hive.log.1), instead of the expected 100 MB file size and 20 backup files.

Supply a custom log4j configuration file with the correct structure by using the ATLAS_OPTS environment variable before running the import script:

```
<?xml version="1.0" encoding="UTF-8" ?>
<!DOCTYPE log4j:configuration SYSTEM "log4j.dtd">
<log4j:configuration xmlns:log4j="http://jakarta.apache.org/log4j/">
  <appender name="console" class="org.apache.log4j.ConsoleAppender">
    <param name="Target" value="System.out"/>
    <layout class="org.apache.log4j.PatternLayout">
      <param name="ConversionPattern" value="%d %-5p - [%t:
%x] ~ %m (%C{1}:%L)%n"/>
    </layout>
  </appender>
  <appender name="FILE" class="org.apache.log4j.RollingFileAppender">
    <param name="File" value="${atlas.log.dir}/${atlas.log.file}"/>
    <param name="Append" value="true"/>
    <param name="maxFileSize" value="100MB"/>
    <param name="maxBackupIndex" value="20"/>
    <layout class="org.apache.log4j.PatternLayout">
      <param name="ConversionPattern" value="%d %-5p - [%t:
%x] ~ %m (%C{1}:%L)%n"/>
    </layout>
  </appender>
  <logger name="org.apache.atlas" additivity="false">
    <level value="info"/>
    <appender-ref ref="FILE"/>
  </logger>
</log4j:configuration>
```

```

    </logger>
    <logger name="org.apache.kafka.common.config.AbstractConfig"
    additivity="false">
      <level value="error"/>
      <appender-ref ref="FILE"/>
    </logger>
  </root>
  <priority value="warn"/>
  <appender-ref ref="FILE"/>
</root>
</log4j:configuration>

```

Export the ATLAS_OPTS variable to point to your custom file before running the import script:

```

export ATLAS_OPTS="-Dlog4j.configuration=file:/path/to/custom-hi
ve-import-log4j.xml"
./import-hive.sh

```

CDPD-87670: [Atlas UI] Apache Atlas Glossary becomes unresponsive when the page size is set to 50

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

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-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-6565: Issue with ddlQueries in Atlas for Hive/Impala tables created in default/non-default db using DAS/HUE respectively

ddlQueries are not created for Impala origin table. ddlQueries created for Impala/Hive CTAS tables have difference in names as compared to same operation for Hive table. The Impala table has the query in the ddlQueries entity name.

None

CDPD-55301: ddlQueries and ALTERNATE_* lineage missing for Spark tables created through spark3-shell

The ddlQueries and ALTERNATE_* lineage missing for Spark tables created using spark3-shell.

None

CDPD-56085: 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

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*.

None

CDPD-58554: Discard audits of specific classification, label, and business metadata

Support to control audits for specific classification, label, and business metadata is not present in the Custom Audit Filters feature.

None

CDPD-40346: The ddlQueries and ALTERNATIVE_ADDCOLS lineage missing for Impala tables.

The ALTERNATIVE_ADDCOLS lineage has some issue when an Impala table is altered and the corresponding lineage is not created.

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

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

None

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

Adding Japanese labels or user defined properties results in the error message: "Invalid label: ###, label should contain alphanumeric characters, _ or -"

None

CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties

Entities, such as hdfs_path, are not returned by the search when their indexable properties are searched using partial search terms made of Japanese or Chinese characters. Only exact matches return results when searching the indexed properties made of Japanese or Chinese characters.

None

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

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

None.

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

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

None.

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

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

None.

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

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

None.

CDPD-43772: Performance issues with Atlas service

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

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

OPSAPS-67783: During rolling upgrade one among two Atlas server failed to start but Cloudera Manager considered as success

Cloudera Manager marks the Execute command Start on service Atlas-1 as a success even when Atlas service had failed to start successfully. In such cases, Atlas logs give the exact reason for Atlas start-up failure.

None

CDPD-19358: "IsIndexable" and "isOptional" value of a typedef's attribute is modified post migration.

In HDP-265, falcon_feed_creation has "stored-in" attribute which has IsIndexable value set to True.

Post migration, "stored-in" is moved to relationshipAttributeDefs and has IsIndexable value set to False.

Similarly, in HDP-265 the table attribute of hive_storage_desc isOptional is set to true. Post migration, isOptional set to False.

None.

CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command

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

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

In converting content from Navigator to Atlas, the create time for Hive tables is not moved to Atlas.

None

CDPD-10574: Suggestion order doesn't match search weights

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

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) does not 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 the metadata collection from HMS. For example, if an ALTER TABLE operation in Spark that is changing a table name and is reporting 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 "@" does not 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, and use 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 does 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 logout 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 CDP.

None.

CDPD-1664: Guest users are redirected incorrectly

Authenticated users logging in to Atlas are redirected to the CDP 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: The 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, to add metadata tags to relationships that you wanted to make sure were unique in the system, you can design the relationship attribute with the property IsUnique equal true. However, in this release, the IsUnique attribute is not enforced.

None.

DOCS-13759: Tag Propagation stops after a certain depth while the lineage is being extended

When a tag is added to an entity at timestamp T1, the entities along the lineage to which the tag must be propagated is calculated at T1. Before tag propagation completes, if the lineage is extended, tag does not propagate to the entities in the extended lineage.

CDPD-41142: When a Kafka console consumer group is run, more than one update audits are seen

After running the console consumer with a consumer group, verify the consumer group entity created, along with the metrics and notifications for the consumer group and topic. The expected result can be: one ENTITY_CREATE audit and one ENTITY_UPDATE audit. But more than one ENTITY_UPDATE audits are seen.

CDPD-40165: Two audits are created for SPARK CTAS table

When following Spark queries are fired:

```
spark.sql("create table table1(id int)")
spark.sql("create table table2 as select * from table1")
```

HMS sends "ENTITY_CREATE" and "ENTITY_FULL_UPDATE_V2".

The extra ENTITY_FULL_UPDATE_V2 message received from HMS is sent as a part of ALTER_TABLE_ADDCOLS event from the HMS Hook side. This behaviour is observed only when the queries are run from Spark SQL and not when run the same queries are run from Beeline.

CDPD-39197: Debug metrics returns empty data

When debug metrics is enabled and some operations are performed, the response is empty

CDPD-36495: Updating legacyAttribute from False to True resets the initially created relationshipAttributes values

Creating types, entities, and to start, you must set the relationship with is_legacy_attribute value as **False**.

Later, update the value relationshipDef is_legacy_attribute to True.

For the entities that were created before updating the is_legacy_attribute to True, the relationship Attributes value is reset.

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-29409: Hive import: Suggestion suggests entity which is deleted.

Suggestions suggests tables of a database, which is a deleted entity.

CDPD-25152: Tag propagation through deferred actions consumes additional time as compared to default flow

The additional time might be due to the small overhead added to create / update task vertex and which is run in the background. This also depends on number of tasks queued to be executed.

CDPD-42954: Zeppelin notebook fails after enabling Atlas-HDFS hook

The Zeppelin notebooks are failing with errors after enabling Atlas-HDFS hook in the CDP cluster.

When the below properties are set for atlas-client.properties in Cloudera Manager:

- atlas.jaas.KafkaClient.option.keyTab
- atlas.jaas.KafkaClient.option.principal

Along with adding the properties in /etc/atlas/conf/atlas-application.properties, Cloudera Manager also adds these properties to atlas-application.properties for other services (like Spark).

Adding these properties interferes with the normal flow of the services (like Spark)

To enable HDFS lineage feature, instead of setting these properties through Cloudera Manager, users can manually add the properties directly in `/etc/atlas/conf/atlas-application.properties`

CDPD-10576: Deleted Business Metadata attributes appear in Search Suggestions

Atlas search suggestions continue to show Business Metadata attributes even if the attributes have been deleted.

None.

CDPD-39427: [HDFS Lineage]: When the input is a directory in case of `put/copyfromLocal/cp/mv` commands, lineage is not created even though the script succeeds.

When Source is a directory and target is a directory which is already present in Atlas, the command succeeds and inserts the data in the desired location, but lineage is not created.

DOCS-13760: System Attributes search, `__classificationNames`: Search with parent tag does not return entities associated to its children tags

System attribute search with `__classificationNames = parent_tag` returns entities associated to `parent_tag` only and not entities associated to its children tag.

Workaround: Instead of using system attribute, employ the basic search attribute "classification" which lists entities associated with inherited classifications.

CDPD-35818: Basic search with tag filter provides approximateCount as -1 when there is no match and is 0 otherwise

When the following search operations are performed:

- Faceted search with both tag filter and entity filter

The observed approximateCount is -1.

Here both entity filter and tag filter are present and when there is no match the response received is -1.

- Faceted search with only entity filter

Performing a basic search provides an approximate value of 0 when there is no match.

- Faceted search with only tag filter

Whenever there is tag filter in the query and there is no entity match, the approximateCount is -1 and if the tag filter is not available, the response approximateCount is 0

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet `atlas.authentication.method.file` is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export

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

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](#)

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

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

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

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

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

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

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

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

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-88230: Bulk import operations currently do not construct lineage as expected

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

Bulk import operations do not create audits for the imported entities.

Known Issues in Cruise Control

Learn about the known issues in Cruise Control, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14**Rebalancing with Cruise Control does not work due to the metric reporter failing to report the CPU usage metric**

On the Kafka broker, the Cruise control metric reporter plugin may fail to report the CPU usage metric.

If the CPU usage metric is not reported, the `numValidWindows` in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.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 the [Moving a Host Between Clusters](#).



Note: Cluster nodes affected by this issue are not displayed as unhealthy.

Known Issues in Hive

Learn about the known issues in Hive, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14**CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress**

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

The following services support DDL statements:

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

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

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

CDPD-60770: Beeline Authentication Issue with Special Characters in Passwords

When LDAP is enabled, users cannot authenticate with Beeline if the password contains a special character. For example, the following string fails:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>;user=user@XXX;password='R3G#xpXyoy1M0Jb1'
```

Use the `-p` parameter to execute the Beeline command:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>; -n user@XXX -p 'R3G#xpXyoy1M0Jb1'
```

CDPD-43769: Inconsistent behavior with CHAR comparisons and string literals

In Hive, 'char' data to 'string' data comparison is considered as equal, but 'char' data to 'varchar' data comparison is considered as not equal if the data has trailing spaces. For example,

```
create table test(colchar char(10), colvarchar varchar(20), colstring string);
insert into test values ('a ', 'a ', 'a ');

select count(*) from test where colchar = colstring;
Output returns 1 because char and string data is considered as equal

select count(*) from test where colchar = colvarchar;
Output returns 0 because char and varchar data is considered as not equal
```

It is recommended that you use 'varchar or string' data type instead of 'char' data type for any character or string representation.

CDPD-42726: Wrong results for agg queries from stats after running multi-insert query on managed table

Multiple inserts at the same time into the same partition results in invalid stats. For example:

```
create table source(p int, key int, value string);
insert into source(p, key, value) values (101,42,'string42');

create table stats_part(key int,value string) partitioned by (p int);
```

```

from source
insert into stats_part select key, value, p
insert into stats_part select key, value, p;

select count(*) from stats_part;
-- In this scenario, StatsOptimizer helps serving this query be
cause the result should be rowNum of the partition p=101.
-- The result is 1, however, the result should be 2.

```

Rewrite the multi-insert statement by merging the branches inserting into the same partition to one branch.

CDPD-57574: Query execution fails due to NullPointerException in DagUtils.setupQuickStart

Hive queries may fail at runtime with a NullPointerException while executing the Tez graph at DagUtils.setupQuickStart method. A part of the stack trace from the error is shown below:

```

ERROR : Failed to execute tez graph.
java.lang.NullPointerException: null at org.apache.hadoop.hive.
ql.exec.tez.DagUtils.setupQuickStart(DagUtils.java:1724)

```

The error usually occurs from an invalid plan that has cycles created by semi-join or map-join edges.

Disable some of the optimizations in order to remove the cycles from the plan. It is sufficient to set one of the following properties:

- set hive.tez.dynamic.partition.pruning=false;
- set hive.tez.dynamic.semijoin.reduction=false;
- set hive.auto.convert.join=false;

To avoid any undesired performance regressions, disable the optimization(s) only for specific queries and not globally (through Cloudera Manager) for the whole workload or cluster.

CDPD-45370: Cartesian join fails with an empty table

When the configuration hive.tez.cartesian-product.enabled is set to true, query fails if one of the tables involved in a Cartesian join has zero rows. This results in an IllegalArgumentException with the message Number of items is 0. Should be positive.

To address this issue, set hive.tez.cartesian-product.enabled to false before running the query.

Apache Jira: [\[HIVE-26582\]](#)

CDPD-43107: SemanticException for INSERT INTO statement when Hive Cost-based Optimizer (CBO) is disabled

If you are running the INSERT INTO query and hive.cbo.enable is set to "false", the query fails with a SemanticException. For example,

```

set hive.cbo.enable=false;

CREATE TABLE mytable (
  id INT,
  str STRING
);

INSERT INTO mytable (id, str) VALUES (1, 'a');

```

```

Output:
org.apache.hadoop.hive.ql.parse.SemanticException: 0:0 Expected
2 columns for insclause-0/default@mytable; select produces 1 col
umns. Error encountered near token 'a'

```

This issue occurs only when columns are specified for the target table in the query. INSERT INTO mytable values (1, 'a'); does not result in an exception.

Enable CBO and run the query.

CDPD-43957: HiveServer shuts down during replication due to high resource usage

During Hive replication (both bootstrap and incremental), you may notice that the HiveServer (HS2) shuts down periodically with the following error:

```
java.sql.SQLException: org.apache.hive.jdbc.ZooKeeperHiveClientException: Unable to read HiveServer2 configs from ZooKeeper
    at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.java:265)
    at org.apache.hive.jdbc.HiveDriver.connect(HiveDriver.java:107)
    at java.sql.DriverManager.getConnection(DriverManager.java:664)
    at java.sql.DriverManager.getConnection(DriverManager.java:247)
    at com.cloudera.enterprise.hive3qt.Hive3QueryTool$HiveOperation.execute(Hive3QueryTool.java:682)
    at com.cloudera.enterprise.hive3qt.Hive3QueryTool.main(Hive3QueryTool.java:935)
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Unable to read HiveServer2 configs from ZooKeeper
    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:177)
    at org.apache.hive.jdbc.Utills.configureConnParamsFromZooKeeper(Utills.java:580)
    at org.apache.hive.jdbc.Utills.parseURL(Utills.java:391)
    at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.java:263)
    ... 5 more
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Tried all existing HiveServer2 uris from ZooKeeper.
    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.getServerHosts(ZooKeeperHiveClientHelper.java:132)
    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:172)
    ... 8 mor
```

This issue occurs when you replicate at scale with an unbalanced cluster setup that has all the roles running on the Cloudera Manager host. As a result, Cloudera Manager ends up in a bottleneck situation because HS2 crashes, preventing further replication. The logs (/var/log/messages) indicate that there were a large number of Java processes running on the host, which exhausted the host's memory and triggered an Out Of Memory Killer process to stop HS2.

You can restart HS2 and reinitiate replication. However, the replication process may take a longer time to complete.

- Follow these recommendations when setting up your source and target clusters:
 - Ensure that no HS2, Hive metastore (HMS), DataNode, or NameNode roles are running on the Cloudera Manager host.
 - Ensure that you have multiple instances of HS2 and HMS roles on different nodes.
 - It is recommended that you have the HS2 and HMS roles on different nodes than the DataNodes or NameNodes.

These practices increase the possibility of Hive replication at scale completing successfully.

- Limit the scale of Hive replication by temporarily disabling replication policies.

CDPD-40730: Parquet change can cause incompatibility

Parquet files written by the parquet-mr library in CDP 7.1.9, where the schema contains a timestamp with no UTC conversion will not be compatible with older versions of Parquet readers. The effect is that the older versions will still consider these timestamps as they would require UTC

conversions and will thus end up with a wrong result. You can encounter this problem only when you write Parquet-based tables using Hive, and tables have the non-default configuration `hive.parquet.write.int64.timestamp=true`.

None.

CDPD-41274: HWC + Oozie issue: Could not open client transport with JDBC Uri

Currently only Spark cluster mode is supported in the Oozie Spark Action with Hive Warehouse Connector (HWC).

Use Spark action in cluster mode.

```
<spark xmlns="uri:oozie:spark-action:1.0">
  ...
  <mode>cluster</mode>
  ...
</spark>
```

CDPD-26556 After an upgrade, querying a CTAS table under certain conditions might throw an exception

If you upgrade your Hive cluster from CDH 6 to CDP 7, create a CTAS table in the CDP cluster from a table you upgraded from CDH, you might see the following exception when you query the new table:

```
class org.apache.hadoop.io.IntWritable cannot be cast to class o
rg.apache.hadoop.hive.serde2.objectinspector.StandardUnionObject
Inspector$StandardUnion
```

This issue involves CDH-based tables having columns of complex types ARRAY, MAP, and STRUCT.

CDPD-23506: OutOfMemoryError in LLAP

Long running spark-shell applications can leave sessions in interactive Hiveserver2 until the Spark application finishes (user exists from spark-shell), causing memory pressure in case of a high number of queries in the same shell (1000+).

You must close spark-shell so that sessions are closed. Add the owner of the database or the tables as a user with read or read/write access to the tables directly.

CDPD-23041: DROP TABLE on a table having an index does not work

If you migrate a Hive table to CDP having an index, DROP TABLE does not drop the table. Hive no longer supports indexes ([HIVE-18448](#)). A foreign key constraint on the indexed table prevents dropping the table. Attempting to drop such a table results in the following error:

```
java.sql.BatchUpdateException: Cannot delete or update a parent
row: a foreign key constraint fails ("hive"."IDXS", CONSTRAINT "
IDXS_FK1" FOREIGN KEY ("ORIG_TBL_ID") REFERENCES "TBLS" ("TBL_ID"
))
```

There are two workarounds:

- Drop the foreign key "IDXS_FK1" on the "IDXS" table within the metastore. You can also manually drop indexes, but do not cascade any drops because the IDXS table includes references to "TBLS".
- Launch an older version of Hive, such as Hive 2.3 that includes IDXS in the DDL, and then drop the indexes as described in [Language Manual Indexing](#).

Apache Issue: [Hive-24815](#)

CDPD-17766: Queries fail when using spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive.

Do not specify `spark.sql.hive.hiveserver2.jdbc.url.principal` in the JDBC URL to invoke Hive remotely.

Workaround: specify `principal=hive.server2.authentication.kerberos.principal` as shown in the following syntax:

```
jdbc:hive://<host>:<port>/<dbName>;principal=hive.server2.authentication.kerberos.principal;<otherSessionConfs>?<hiveConfs>#<hiveVars>
```

CDPD-13636: Hive job fails with OutOfMemory exception in the Azure DE cluster

Set the parameter `hive.optimize.sort.dynamic.partition.threshold=0`. Add this parameter in Cloudera Manager (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)

CDPD-10848: HiveServer Web UI displays incorrect data

If you enabled auto-TLS for TLS encryption, the HiveServer2 Web UI does not display the correct data in the following tables: Active Sessions, Open Queries, Last Max n Closed Queries

Known Issues in HDFS

Learn about the known issues in HDFS, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

There are no known issues identified in the releases before Cloudera Runtime 7.1.9 SP1 CHF 14.

Known Issues in Hue

Learn about the known issues in Hue, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14**CDPD-94739: Hue audit logging requires debug mode**

In Hue, audit logging is only active when debug logging is enabled by setting `DJANGO_DEBUG_MODE = true`. By default, audit logging is disabled.

Update to the upcoming release to restore audit logging functionality.

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

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark queries in a batch mode launches a shell script Oozie job from Hue and this is not supported on a secure cluster.

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

OPSAPS-69659: Hue service fails on restart with "Unable to find pycpg2 2.5.4" error

Hue service fails to restart and you see the following error: Unable to find pycpg2 2.5.4. This could be because you have installed Python in a non-default location and Hue is unable to locate the pycpg2 PostgreSQL database adapter.

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

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

Key: PYTHONPATH

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

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

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

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

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

None.

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

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

None.

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

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

None.

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

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

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

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

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

CDPD-43293: Unable to import Impala table using Importer

Creating Impala tables using the Hue Importer may fail.

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

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

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

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

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

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

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

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

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

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

Hue 4 web interface uses the unsafe-inline directive in its CSP header. As a result, the application server does not set the CSP header in its HTTP responses, and therefore does not benefit from the additional protection against potential cross-site scripting issues and other modern application vulnerabilities which a properly configured CSP may provide. This could lead to application vulnerability.

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

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

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

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

```
[search]
# URL of the Solr Server
solr_url=http://[***HOSTNAME***]:[***PORT***]/solr/
```

For example:

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

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

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

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

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

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

Invalid S3 URI error while accessing S3 bucket

The Hue Load Balancer merges the double slashes (//) in the S3 URI into a single slash (/) so that the URI prefix "/filebrowser/view=S3A:/" is changed to "/filebrowser/view=S3A:". This results in an error when you try to access the S3 buckets from the Hue File Browser through the port 8889.

The Hue web UI displays the following error: “Unknown error occurred”.

The Hue server logs record the “ValueError: Invalid S3 URI: S3A” error.

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.

You should be able to load the S3 browser from both 8888 and 8889 ports.

Alternatively, you can use the Hue server port 8888 instead of the load balancer port 8889 to resolve this issue.

Error while rerunning Oozie workflow

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

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

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

CDPD-17956, CDPD-14167: Setting idle session timeout for Hue does not work when the cluster is secured using Knox SSO

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

None

Query history displays "Result Expired" message

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

None.

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

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

Hue limitation after upgrading from CDH to CDP Private Cloud Base

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

Hue UI is blank upon login after upgrading to CDP 7.1.7 from CDH 6

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

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

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



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

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

Unsupported features

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

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

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

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

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

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

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

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

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

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



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

Known Issues in Impala

Learn about the known issues in Impala, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

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

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

The following services support DDL statements:

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

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

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

CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the `--tls_ciphersuites` startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.

Ensure that the list in the `--tls_ciphersuites` startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

Apache Jira: [IMPALA-14625](#)

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work

If you try to run `impala-shell` on RHEL 9 by setting the default python executable available in `PATH` to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, `IMPALA_PYTHON_EXECUTABLE=python3`.

CDPD-81980: Impala query may fail due to backpressure

Impala queries could fail with a dropped due to backpressure error, even when service queue and memory usage appear normal.

Increase the value of the `--mem_limit` startup option.

CDPD-42958: After upgrading the CDH 7.1.9 from CDH 6.x, under certain conditions you cannot insert data into a table

Under the following conditions, after upgrading from CDH 6.x to CDH 7.1.9 you cannot insert data into a table from Impala:

- On CDH 6.x, you created a database with Impala in a user specified HDFS location.
- Using Hive, you then created a table in the database.

Under these conditions, the database and table are stored in the user-specified HDFS directory. After upgrading, the HDFS directory of the table is read-only for Impala. Consequently, from Impala you cannot insert new data into the table because Impala does not have write permission on the HDFS directory.

Workaround: To resolve this issue, use either one of the following workarounds:

- Using the Ranger Web UI, in the policy repository `cm_hdfs`, grant the user 'impala' write permission on the directory where the table resides.
- Enter the following command to grant write permission to user 'impala' on the HDFS directory where the table resides.

```
hdfs dfs -setfacl -m default:user:impala:rwX <HDFS directory>
```

CDPD-80019: Impala fails to start due to large catalog-update topic

Impala coordinators fail to start if the catalog-update statestore topic exceeds 2GB, hitting the Thrift 0.16.0 message size limit (`thrift_rpc_max_message_size`). This prevents the coordinator from processing the update.

- Enable local catalog mode: `--use_local_catalog=true`
- In Cloudera Manager, go to Impala Configuration search for "Catalog Server Command Line Argument Advanced Configuration Snippet (Safety Valve)" and click 
- Add property name and value:

```
Name : --catalog_topic_mode
Value: minimal
```

- Save changes.

Apache Jira: [IMPALA-12350](#)

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/<data
base>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

You must use a single Knox node to access Impala UI.

Impala api calls via 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-21828: Multiple permission assignment through grant is not working

None

Problem configuring masking on tables using Ranger

The following Knowledge Base article describes the behavior when we configure masking on tables using Ranger. This configuration works for Hive, but breaks queries in some scenarios for Impala.

For a workaround, see the following Knowledge Base article: [ERROR: "AnalysisException: No matching function with signature: mask\(FLOAT\)" when Impala jobs fail with the following error with signature: mask\(FLOAT\)](#)

IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXXXX"

IMPALA-532: Impala should tolerate bad locale settings

If the LC_* environment variables specify an unsupported locale, Impala does not start.

Add LC_ALL="C" to the environment settings for both the Impala daemon and the Statestore daemon.

Avro Scanner fails to parse some schemas

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

IMPALA-691: Process mem limit does not account for the JVM's memory usage

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the impalad daemon.

To monitor overall memory usage, use the top command, or add the memory figures in the Impala web UI /memz tab to JVM memory usage shown on the /metrics tab.

IMPALA-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Remove trailing semicolon from the Avro schema.

IMPALA-1652: Incorrect results with basic predicate on CHAR typed column

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the RPAD() function to blank-pad literals compared with CHAR columns to the expected length.

IMPALA-1821: Casting scenarios with invalid/inconsistent results

Using a CAST() function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

IMPALA-2005: A failed CTAS does not drop the table if the insert fails

If a CREATE TABLE AS SELECT operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

Drop the new table manually after a failed CREATE TABLE AS SELECT

IMPALA-2422: % escaping does not work correctly when occurs at the end in a LIKE clause

If the final character in the RHS argument of a LIKE operator is an escaped \% character, it does not match a % final character of the LHS argument.

None

IMPALA-2603: Crash: impala::Coordinator::ValidateCollectionSlots

A query could encounter a serious error if includes multiple nested levels of INNER JOIN clauses involving subqueries.

None

IMPALA-3094: Incorrect result due to constant evaluation in query with outer join

An OUTER JOIN query could omit some expected result rows due to a constant such as FALSE in another join clause. For example:

```
explain SELECT 1 FROM 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
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

IMPALA-3509: Breakpad minidumps can be very large when the thread count is high

The size of the breakpad minidump files grows linearly with the number of threads. By default, each thread adds 8 KB to the minidump size. Minidump files could consume significant disk space when the daemons have a high number of threads.

Add `-\-minidump_size_limit_hint_kb=size` to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then

2 KB per thread after that.) The minidump file can still grow larger than the "hinted" size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

IMPALA-4978: Impala requires FQDN from hostname command on Kerberized clusters

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the `gethostname()` system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a Kerberized cluster.

Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname` only returns the short name, pass the command-line flag `##hostname=FULLY_QUALIFIED_DOMAIN_NAME` in the startup options of all Impala-related daemons.

IMPALA-6671: Metadata operations block read-only operations on unrelated tables

Metadata operations that change the state of a table, like `COMPUTE STATS` or `ALTER RECOVER PARTITIONS`, may delay metadata propagation of unrelated unloaded tables triggered by statements like `DESCRIBE` or `SELECT` queries.

None

IMPALA-7072: Impala does not support Heimdal Kerberos

None

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.

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.

Known Issues in Iceberg

Learn about the known issues in Iceberg, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

There are no known issues identified in earlier releases.

Known Issues in Kafka

Learn about the known issues in Kafka, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

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

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

The following services support DDL statements:

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

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

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

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker) in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Save Changes > Restart SMM.

The offsets.topic.replication.factor property must be less than or equal to the number of live brokers

The offsets.topic.replication.factor broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a GROUP_COORDINATOR_NOT_AVAILABLE error until the cluster size meets this replication factor requirement.

None

Requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true

The first few produce requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true.

Increase the number of retries in the producer configuration setting retries.

Performance degradation when SSL Is enabled

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with ssl.secure.random.implementation = SHA1PRNG. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

OPSAPS-43236: Kafka garbage collection logs are written to the process directory

By default Kafka garbage collection logs are written to the agent process directory. Changing the default path for these log files is currently unsupported.

None

RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transactional method because we are in an error state
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeFailWithError(TransactionManager.java:1125)
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeAddPartition(TransactionManager.java:442)
    at org.apache.kafka.clients.producer.KafkaProducer.doSend(KafkaProducer.java:1000)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:914)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:800)
    .
    .
    .
    Caused by: org.apache.kafka.common.errors.ClusterAuthorizationException: Cluster authorization failed.
```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting `enable.idempotence` to `false`.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

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

The Debezium Db2 Source connector does not support the evolution (updates) of schemas. In addition, schema change events are not emitted to the schema change topic if there is a change in the schema of a table that is in capture mode. For more information, see [DBZ-4990](#).

None.

CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy compression. The HDFS Stateless Sink connector is only affected if the Compression Codec or Compression Codec for Parquet properties are set to SNAPPY.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: org.apache.hadoop.util.NativeCodeLoader.buildSupportsSnappy()
```

```
Failed to write to HDFS due to java.lang.RuntimeException: native snappy library not available: this version of libhadoop was built without snappy support.
```

Download and deploy missing libraries.



Important: Ensure that you complete steps 1-11 on all Kafka Connect hosts. Additionally, ensure that the advanced configuration snippet in step 12 is configured for all Kafka Connect role instances.

1. Create the `/opt/nativelibs` directory.

```
mkdir /opt/nativelibs
```

2. Change the owner to kafka.

```
chown kafka:kafka /opt/nativelibs
```

3. Locate the directory containing the Hadoop native libraries and copy its contents to the directory you created.

```
cp /opt/cloudera/parcels/CDH/lib/hadoop/lib/native/* /opt/nativelibs
```

4. Verify that `libsnapy.so` was copied to the directory you created.
5. Remove the following from `/opt/nativelibs`.

```
libhadoop.a
libhadoop.so
libhadoop.so.1.0.0
```

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

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Unsupported Features

The following Kafka features are not supported in Cloudera Data Platform:

- Only Java and .Net based clients are supported. Clients developed with C, C++, Python, and other languages are currently not supported.
- The Kafka default authorizer is not supported. This includes setting ACLs and all related APIs, broker functionality, and command-line tools.
- SASL/SCRAM is only supported for delegation token based authentication. It is not supported as a standalone authentication mechanism.
- Kafka KRaft in this release of Cloudera Runtime is in technical preview and does not support the following:
 - Deployments with multiple log directories. This includes deployments that use JBOD for storage.
 - Delegation token based authentication.
 - Migrating an already running Kafka service from ZooKeeper to KRaft.
 - Atlas Integration.

Limitations

Collection of Partition Level Metrics May Cause Cloudera Manager's Performance to Degrade

If the Kafka service operates with a large number of partitions, collection of partition level metrics may cause Cloudera Manager's performance to degrade.

If you are observing performance degradation and your cluster is operating with a high number of partitions, you can choose to disable the collection of partition level metrics.



Important: If you are using SMM to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both SMM and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, SMM will not be able to display information about partitions. In addition, Cruise Control will not operate properly.

Complete the following steps to turn off the collection of partition level metrics:

1. Obtain the Kafka service name:
 - a. In Cloudera Manager, Select the Kafka service.
 - b. Select any available chart, and select Open in Chart Builder from the configuration icon drop-down.
 - c. Find \$SERVICENAME= near the top of the display.

The Kafka service name is the value of \$SERVICENAME.

2. Turn off the collection of partition level metrics:
 - a. Go to Hosts Hosts Configuration .
 - b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

Enter the following to turn off the collection of partition level metrics:

```
[KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_entity_update_enabled=false
```

Replace [KAFKA_SERVICE_NAME] with the service name of Kafka obtained in step 1. The service name should always be in lower case.

- c. Click Save Changes.

Known Issues in Knox

Learn about the known issues in Knox, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

CDPD-92941: Knox Gateway logging configuration cannot be customized through Cloudera Manager

The Knox Gateway Logging Configuration (Safety Valve) in Cloudera Manager is not applied to Knox Gateway. Individual logger configurations cannot be set or modified through the Cloudera Manager UI.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

CDPD-76435: Default HTTP client timeout for Streams Messaging Manager requests may cause connection timeouts

Some Streams Messaging Manager requests, particularly those retrieving topic data, may take up to 30 seconds to complete, even when there are no connectivity issues. Since the default gateway timeout is set to 20 seconds, this can result in connection timeouts despite the system functioning as expected.

Increase the HTTP client and socket timeouts for Streams Messaging Manager services to 60 seconds. This prevents connection timeouts when the requests take longer than the default gateway timeout.

Follow these steps to configure the required properties for the SMM API and SMM UI services:

1. In Cloudera Manager, select the Knox service.
2. Go to Configuration.
3. Find the Knox Simplified Topology Management - cdp-proxy property, and click  to add the following lines:

```
SMM-UI:replayBufferSize=512
SMM-UI:httpclient.connectionTimeout=1m
SMM-UI:client.socketTimeout=1m
SMM-API:replayBufferSize=512
```

```
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

4. Find the Knox Simplified Topology Management - cdp-proxy-api property, and click **+** to add the following lines:

```
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

5. Find the Knox Simplified Topology Management - cdp-proxy-token property, and click **+** to add the following lines:

```
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

6. Save your changes.
7. Refresh the cluster.

Known Issues in Kudu

Learn about the known issues in Kudu, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

CDPD-94474: Kudu Tablet Server Crash on RHEL9

Kudu tablet servers kudu-tserver crash on Red Hat Enterprise Linux 9 (RHEL9) x86_64 platforms when the libgcc package is updated to version 11.5.0-11 or newer. This crash manifests as a Segmentation Fault (SIGSEGV) or an Abort (SIGABRT) error under certain workloads.

1. Log in to Cloudera Manager as an Administrator.
2. Navigate to Cloudera Manager Clusters Kudu Configuration .
3. Navigate to the Configuration tab, and add the following value in the Tablet Server Advanced Configuration Snippet (Safety Valve) for gflagfile for the Tablet Server Default Group field:

```
--mrs_use_codegen=false
```

4. Click Save Changes.
5. Restart the Kudu Service.

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.

Workaround: Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

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.

Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

Unsupported features

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 in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

No known issues identified in Livy.

Known issues in Navigator Encrypt

Learn about the known issues in Navigator Encrypt, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

Failed to start navencrypt-mount.service

After installing NavEncrypt on SLES, if the command "systemctl status navencrypt-mount" fails with the error: "Failed to start navencrypt-mount.service: Unit navencrypt-mount.service failed to load: No such file or directory", the systemd files need to be reloaded.

Run the command:

```
(sudo) systemctl daemon-reload
```

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

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 Oozie

Learn about the known issues in Oozie, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

There are no known issues identified in the releases before Cloudera Runtime 7.1.9 SP1 CHF 14.

Known Issues in Ozone

Learn about the known issues in Ozone, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

CDPD-88335: Datanode decommission fails when other datanodes are offline due to invalid affinity node in Ratis replication

Datanode decommissioning fails if one or more datanodes holding Ratis replicas are offline. During decommission, the Storage Container Manager (SCM) tries to replicate Ratis container data to maintain redundancy. However, SCM's placement logic still references an affinity node that was previously part of the network topology but has since been removed due to the datanode going offline. This causes repeated placement errors and prevents the decommission process from completing successfully.

Apache Issue: [HDDS-13544](#)

CDPD-87429: Unexpected log lines in Ozone snapshot diff output

`ozone.client.server-defaults.validity.period.ms` configuration introduced to get server default configurations causes Replication Manager to fail when replicating data between Ozone buckets with the following error when the config value does not specify a timeunit.

```
...
java.lang.IllegalStateException: Invalid first line in diff: "25
/07/02 05:58:44 INFO Configuration.deprecation: No unit for ozon
e.client.server-defaults.validity.period.ms(3600000) assuming MI
LLISECONDS"
    at com.cloudera.enterprise.distcp.OzoneCopyListingUtils.parseOz
oneDiff(OzoneCopyListingUtils.java:397)
...
```

Overwrite the `ozone.client.server-defaults.validity.period.ms` to configure the value to have time unit (ns, ms, s, m, h). For example, 3600000ms.

CDPD-71648: Snapshot creation failing with "PERMISSION_DENIED" error on a bucket created during pre-upgrade in 7.1.7 SP3

On upgrade of cluster from 7.1.7 SP3 to 7.1.9 SP1, we are not able to create snapshots on a pre-upgrade volume or bucket with "PERMISSION_DENIED" error, since in 7.1.7SP3, we were using full KRB owner name with domain but currently, we only use and search for user shortname.

Update the owner of the bucket with the user shortname and then create snapshot by running the commands:

```
ozone sh bucket update --user=[{***USER_SHORT_NAME***}] {volume}/
{bucket}
ozone sh snapshot create {volume}/{bucket}
[{***SNAPSHOT_NAME***}]
```

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

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

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

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

```

</property>

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

```

CDPD-64865: Intermittent timeout in TestBlockDeletion.testBlockDeletion (BlockDeletionService stuck)

If the event of SCM exiting safe mode was triggered multiple times, this might eventually cause the SCMBlockDeletingService to transition into the PAUSING state.

Check the SCM logs and ensure that the SCMBlockDeletingService is not running, and then restart the SCM.

CDPD-81569: Recon is not reflecting correct metrics for Volumes, Buckets, Keys, and Pending Key Deletions

Recon is displaying data for container and space utilisation only. The metrics for Volumes, Buckets, Keys, and Pending Key Deletions are not being updated and displayed correctly.

Apache Jira: [HDDS-12607](#)

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: 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.

The Prometheus binary is not available in CDP Private Cloud Base 7.1.9 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.

After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and if Ozone is in the Non-HA environment, an exception message is observed during the finalization of the Ozone upgrade.

During the finalization of the upgrade, a ClassNotFoundException message for org.cloudera.log4j.redactor.RedactorAppender class may be displayed. The error message can be ignored because the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

CDPD-68951: In CDP Private Cloud Base 7.1.9 CHF2 version and lower, the command `ozone sh key list <bucket_path>` displays the isFile flag in a key's metadata as false even when the key is a file.

This issue is rectified in CDP Private Cloud Base 7.1.9 CHF3. However, the pre-existing (pre-upgrade) key's metadata cannot be changed.

None

CDPD-67771: 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.

CDPD-64398: In SCM with non-HA configuration, Secret key manager does not initialise. Hence, the startup of OM and DataNode fails because it cannot get a secret key. This key is used when security is enabled and used in Block token and container token verification while communicating between the Ozone client, OM, and Datanode.

This issue does not occur in SCM with HA configuration.

You must force exit from the safe mode for SCM. This triggers the initialization of the Secret key manager. Below are the options:

- exit safe mode manually whenever SCM is started
- disable safe mode by setting `hdds.scm.safemode.enabled=false` in safety valve for SCM configuration

Impact of disabling safe mode and the purpose of safe mode:

- Write should not fail once the cluster is out of safe mode
- Read of existing data should not fail after the cluster is out of safe mode
- Unnecessary re-replication should be avoided during cluster restart

So disabling safe mode does not have any major impact on Ozone.

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.

CDPD-52412: keyManagerImpl#listStatus exceeds the maximum RPC length.

To reduce the payload, you must increase the part size.

OPSAPS-68159: If you did not deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Error when distributing parcel to host, Ignoring non-compliant parcel manifest error is displayed after Cloudera Manager is upgraded to 7.11.3.

If you encounter the error, perform the following steps:

1. Deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.11.3.
2. Restart the cluster with a delay of 10 minutes.
3. Continue to upgrade CDH 7.1.8 to CDH 7.1.9.

OPSAPS-68159: If you did not deactivate the Ozone parcel 718.2.x on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Ozone roles during the CDH 7.1.8 upgrade to CDH 7.1.9.

If you encounter the error, perform the following steps:

1. Deactivate the Ozone parcel 718.2.x.
2. Restart the Ozone service.
3. Perform Finalize Upgrade for the Ozone service.

Step result: The Ozone roles will be displayed in green.

CDPD-60989: The packaging version for Apache Ozone points to the 1.2.0 older version. This is a version string problem and not a packaging issue. The version of the Apache Ozone binary is closest to 1.4.0.

None. This only affects the JAR versioning.

CDPD-60366: Native library loader fails when system property `native.lib.tmp.dir` is not set. It fails because the library is copied to `/` instead of the `cwd`.

Setting `native.lib.tmp.dir` to a certain path like `/tmp` should solve the issue. To set this system property, you must add `-Dnative.lib.tmp.dir=/tmp` to `ozone_java_opts` in Cloudera Manager configuration of the Ozone cluster.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-60598: DataNode can consume up to the JVM heap size worth of direct memory buffers under specific failure scenarios. This can lead to a DataNode crash. This is due to Netty allocating and not freeing the direct memory buffers.

Restarting DataNode will resolve this issue.

CDPD-60466: The Ozone client is missing in the Cloudera Manager's Service Monitor node. This causes Ozone Canary's health check to fail.

Ensure you install Ozone CLI on the same node as Cloudera Manager's Service Monitor.

CDPD-60012: The Ozone SCM may be stuck in safe mode after upgrade to 7.1.9. This is due to a bug in how the SCM accounts for open storage containers.

The workaround is to temporarily set the Ozone configuration `hdds.scm.safemode.threshold.pct` to a lower value like 0.90 and restart the SCM.

CDPD-59679: In some scenarios, excess UNHEALTHY replicas of EC containers may not be removed from the cluster.

None

CDPD-50116: Topology-aware reads can provide better performance by routing applications to the closest replica of a file, if available, on the same physical node or same rack. However this feature is disabled by default in CDP 7.1.9.

This feature can be enabled by setting `ozone.network.topology.aware.read` to true using Cloudera Manager safety valve.

CDPD-57165: The DataNode disk usage thread may abort due to Ratis log tailing.

Use `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory` to calculate disk usage. Configure in the `ozone-site.xml` property: `hdds.datanode.du.factory` value: `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory`

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.

CDPD-55513: Hive external table replication policy does not migrate Hive tables in HDFS storage to Ozone.

To replicate data from HDFS to Ozone, see [Migrating your data from HDFS to Ozone](#). You can use HMS Mirror to replicate the metadata.

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"`

OPSAPS-67650: Ozone uses RocksDB as a library to persist metadata locally.

By default, RocksDB places certain executables in `/tmp`, and thus encounters errors when `/tmp` is mounted with `noexec`.

The workaround is to configure RocksDB to put executables at another location. On a PhatCat node, the steps are:

1. Go to Cloudera Manager UI OZONE Configuration .
2. Find Ozone Service Environment Advanced Configuration Snippet (Safety Valve) and set the following environment variable: `ROCKSDB_SHARED_LIB_DIR=/VAR/TMP`

3. Restart Ozone.

CDPD-49137: Ozone Manager Kerberos token expires for SCM communication and OM does not log in again.

Sometimes, OM's Kerberos token is not updated and it stops to communicate with SCM. When this occurs, writes start failing.

Restart OM or set the safety valve `hadoop.kerberos.keytab.login.autorenewal.enabled = true`

CDPD-56684: Keys get deleted when you do not have permission on volume

When a volume is deleted, it recursively deletes the buckets and keys inside it and only then deletes the volume. The volume delete ACL check is done only in the end, due to which you may end up deleting all the data inside the volume without having delete permission on the volume.



Note: There was no bucket/key permission available which allowed the user to delete them recursively.

CDPD-50610: Large file uploads are slow with OPEN and stream data approach

Hue file browser uses the append operation for large files. This API is not supported by Ozone in 7.1.9 and therefore large file uploads can be slow or timeout from the browser.

Use native Ozone client to upload large files instead of the Hue file browser.

OPSAPS-64097: Ozone service restart failed at SCM

Stopping SCM service using Cloudera Manager can sometimes timeout and need a retry. The Cloudera Manager API waits for 90 seconds which is not sufficient under certain circumstances.

Retry the shutdown using Cloudera Manager if the SCM still shows as running after a refresh of the service status.

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-filesystem-hadoop3.jar`, it is likely due to the presence of the user `hdfs` on the node prior to CDP parcel activation. As a result, the Cloudera Manager agent skips the initialization related to HDFS, leading to the non-creation of the `/var/lib/hadoop-hdfs` directory.

Create the directory `“/var/lib/hadoop-hdfs”` on all nodes followed by the deactivation and activation of the CDP parcel (deactivate and activate the Ozone parcel instead, in case Ozone parcel is used).

CDPD-50447: When SCM High Availability is enabled, each of the SCM web UIs report the host of the web UI as the leader of HA, and the other two as followers. This gives wrong information

Correct output is available by running the `ozone admin scm roles --service-id=<ID>` command.

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](#).

OPSAPS-62327: In an Ozone cluster without any gateway roles, Ozone is unable to deploy client configurations and displays the ConfigGenException error.

You must add the Ozone gateway roles to the cluster.

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-36389: The configurations `datanodes.involved.max.percentage.per.iteration` and `size.moved.max.per.iteration` are meant to limit the max number of DataNodes that is involved and maximum size that can move in an iteration. This bug causes balancer to stop an iteration when it is two DataNodes or 1 Container size (5GB) away from hitting these limits. However, these DataNodes can again be considered for balancing in the next iteration. This means the cluster is balanced after enough iterations, albeit a bit slowly. This bug is apparent in small clusters of around four DataNodes where the DataNodes can be either the source or target for a lot of moves but the iteration gets stopped when three DataNodes have been involved. It takes a higher number of iterations to eventually balance this cluster. While this is a performance issue, it does not prevent balancer from ultimately balancing the cluster.

Increase the speed for balancing by decreasing the interval between each iteration using the configuration `balancing.iteration.interval`. The value of this configuration must be greater than `hdds.datanode.du.refresh.period`. `size.moved.max.per.iteration` can be increased to allow more data to move in one iteration.

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-35141: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2> (state=08S01,code=40000) java.sql.SQLException: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2>. This issue does not occur if the source and target buckets are different in Hive queries. Currently, copying across the same bucket is only supported.

Avoid different buckets in source and target path. Ensure that `hive.exec.stagingdir` is unset or set to the same bucket as the warehouse so that staging directories are created in the same bucket as the warehouse.

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-41184: With Legacy buckets, FileSystem Optimised is not interoperating with the Ozone shell command. The cause is that the directory key entry in the DB KeyTable is stored as "dir1/" with trailing slash. But while performing the described operation, Ozone shell (o3://) is normalizing the given path and removing the trailing slash "/" from it. That results in KEY_NOT_FOUND exception.

There are three workarounds:

- Use FileSystem API to Delete the Directories rather than Shell-Command API.
- Use FSO buckets instead of Legacy Buckets. As in FSO, you can create Intermediate Directories and Delete Directories using the Ozone shell commands.
- Disable and set the configuration `ozone.om.enable.filesystem.pathsflag` to false in order to delete the directories. This is generally not a preferred workaround because the cluster must be restarted to pick up the new changes.

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-12542: Mount of Ozone filesystem with the help of FUSE fails.

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-42897: EC writes are failing with "No enough datanodes to choose" after EC replication config is set globally.

EC writes start failing when large number of pipelines are created as a result of multiple EC configs with different chunk sizes used to write keys.

If standard EC configs (that is., rs-3-2-1024k) are used to write keys, number of pipelines created per DataNode is limited to five and this issue is not seen with standard EC configs.

The recommendation is not to create too many random chunk sizes. It is configurable because, users can decide based on their workload. But do not to have separate chunk sizes for each file.

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-40560: Filesystem Operations using Hadoop S3a connector on a FILE_SYSTEM_OPTIMIZED bucket fails. `org.apache.hadoop.ozone.om.exceptions.OMException: Unable to get file status: volume: s3v bucket: fso key: test/`

Do not run Hadoop S3a commands on a FILE_SYSTEM_OPTIMIZED bucket. Use OBJECT_STORE bucket layouts.

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.

CDPD-43432: Ozone Service in fault state in DataNode - Long Running setup.

Upgrade RocksDB to the latest version.

OPSAPS-63999: In the newly installed cluster, the Finish upgrade option is clickable.

None

OPSAPS-64648: Failed to start Ozone node using Cloudera Manager if default log path /var/log/hadoop-ozone does not exist. If this path does not exist, any Ozone nodes(for example SCM or DataNode) restart will fail.

Run the following command `sudo -u hdfs mkdir -p /var/log/hadoop-ozone` or replace hdfs with the user Ozone roles that are running.

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.

Known Issues in Ranger

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

Access issues with users having similar name as admin or keyadmin users

Users with USER role that have a similar name as admin or keyadmin users are able to query those admin and keyadmin users.

None.

UnsupportedClassVersionError

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.

CDPD-68806: The Revoke operation for users belonging to a group or role permission does not function as expected.

List command is listing all the tables even when the user permission is revoked. And also the command does not add any deny policy to Ranger for that specific user.

For example:

user_1 is the part of group_1

user_2 is the part of role_1

Create a Ranger policy for group_1 and role_1 to access all the resources in HBase.

Try to revoke the permission for user_1 and user_2 using HBase shell global revoke command:

```
revoke 'user_1'
```

```
revoke 'user_2'
```

But there is no deny policy added for that user. And access is still granted to user_1 and user_2.

This behavior is currently not supported in HBase shell. Must be handled manually using the Ranger policy change.

CDPD-68739: The revoke command does not work when using the HBase shell

While using the HBase shell, running the revoke command does not cancel the user permission. Users are able to perform actions even after running the revoke command.

For example:

When a user is granted permission to a particular table using the HBase grant command.

```
hbase:003:0> grant 'user_1', 'RWCA', 'some_abc_table'
```

And then execute the global revoke command for that user:

```
hbase:005:0> revoke 'user_1'
```

Permission for that user 'user_1' for the table 'some_abc_table' not got revoked.

None.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-58860: After upgrading CDP-7.1.8 to CDP-7.1.9, cdp-proxy token missing from Knox-Ranger policy.

As part of [OPSAPS-67480](#) in 719 default ranger policy is added from cdp-proxy-token topology, so that after a new installation of CDP-7.1.9, the Knox-Ranger policy includes cdp-proxy-token. However, upgrades do not add cdp-proxy-token to cm_knox policies automatically.

Manually add cdp-proxy-token to the Knox policy, using Ranger Admin Web UI.

1. Log in to Cloudera Manager Ranger Admin Web UI, as a Ranger administrator.
2. On Ranger Admin Web UI Service Manager Resource Knox, click cm_knox.
3. In Knox Policies, open the CDP Proxy UI, API and Token policy.
4. In Knox Topology*, add cdp-proxy-token.
5. Click Save.
6. Restart Ranger.

CDPD-58704: Hadoop roll key / key delete command shows operation failed error when one KMS host is down, even when operation succeeds

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered KMS instances. If one KMS instance is stopped (not deleted), client gets 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.

```
hbase:001:0> revoke 'hrt_11'
ERROR: org.apache.hadoop.hbase.coprocessor.CoprocessorException: HTTP 400 Error: processSecureRevokeRequest processing failed
```

```

        at org.apache.ranger.authorization.hbase.RangerAuthoriza
tionCoproprocessor.preRevoke(RangerAuthorizationCoproprocessor.java:13
09)
        at org.apache.ranger.authorization.hbase.RangerAuthorizat
ionCoproprocessor.preRevoke(RangerAuthorizationCoproprocessor.java:112
8)
        at org.apache.hadoop.hbase.master.MasterCoproprocessorHost$16
2.call(MasterCoproprocessorHost.java:1857)
        at org.apache.hadoop.hbase.master.MasterCoproprocessorHost$1
62.call(MasterCoproprocessorHost.java:1854)
        at org.apache.hadoop.hbase.coprocessor.CoproprocessorHost$0
bserverOperationWithoutResult.callObserver(CoproprocessorHost.java:
558)
        at org.apache.hadoop.hbase.coprocessor.CoproprocessorHost.e
xecOperation(CoproprocessorHost.java:631)
        at org.apache.hadoop.hbase.master.MasterCoproprocessorHost.pr
eRevoke(MasterCoproprocessorHost.java:1854)
        at org.apache.hadoop.hbase.master.MasterRpcServices.revo
ke(MasterRpcServices.java:2740)
        at org.apache.hadoop.hbase.shaded.protobuf.generated.Mast
erProtos$MasterService$2.callBlockingMethod(MasterProtos.java)
        at org.apache.hadoop.hbase.ipc.RpcServer.call(RpcServer.ja
va:387)
        at org.apache.hadoop.hbase.ipc.CallRunner.run(CallRunner.
java:139)
        at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(
RpcExecutor.java:369)
        at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(
RpcExecutor.java:349)

```

None

CDPD-56741: Improvement in log message when jwtauth not used

The following exception is printed at startup only and it is not cluttering logs:

```

2023-05-30 06:18:40,127 ERROR org.apache.ranger.rms.security.RMS
JwtAuthFilter:
    quasar-pibgzl-1.quasar-pibgzl.root.hwx.site-startStop-1]:
    Failed to initialize Ranger RMS JWT Auth Filter.
    java.lang.Exception: RangerJwtAuthHandler:
    Mandatory configs ('jwks.provider-url' & 'jwt.public-key')
    are missing, must provide atleast one.
    at org.apache.ranger.authz.handler.jwt.RangerJwtAuthHan
dler.initialize(RangerJwtAuthHandler.java:84)
    ~[ranger-authn-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
    at org.apache.ranger.rms.security.RMSJwtAuthFilter.init
ialize(RMSJwtAuthFilter.java:77)
    ~[ranger-rms-common-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0
-186]
    at jdk.internal.reflect.NativeMethodAccessorImpl.invoke
0(Native Method) ~[?:?]
    at jdk.internal.reflect.NativeMethodAccessorImpl.invoke(N
ativeMethodAccessorImp

```

This exception is logged, since mandatory JWT auth filter configs are not provided (expected for some environments like y-cloud) Even though JWT auth filter is failed to initialised, it falls back to kerberos auth, hence no impact from auth perspective

NA

CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup

This is a warning log printed in catalina.out file when Ranger RMS server is initialised. Following exception is observed only in Oracle 19 setup. FileNotFoundException: /usr/share/java/oraclepki.jar
NA

CDPD-48975: Ranger KMS KTS to KMS DB migration : keys with the same name but different case are not migrated

KMS keys are not case sensitive

No workarounds. Such keys combination are very rare and migration doc was updated to check such keys before starting the migration.

CDPD-42598: Kafka policy creation allowed with incorrect permissions.

When creating a Kafka policy from the UI, the permissions "Idempotent write" and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource

None.

CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload:

```
{"resourceName": "classification", "userInput": "", "resources":  {"classification": []}}
```

expectation is to return all the classifications . But the response has only "classification" Happens similarly for entity-label , entity-business-metadata.

None.

CDPD-40734: User allowed to insert data into a hive table when there is a deny policy on a table column.

A user is allowed to enter data into a table even if there is a deny policy present on one of the table columns.

Test scenario details:

```
Policy setup :-
  policy 1 :- all access policy for hrt_qa, hive and impala
  users
    resources - database - * , table - *, column - *
    users : hrt_qa, hive, impala
    access - all access allowed
  policy 2 :- policy on test_1.table_1 for hrt_5
  users : hrt_5
    resources : database - test_1, table - table_1, column - *
    access :- all access allowed
  policy 3 :- deny policy on test_1.table_1.c0 for hrt_5
  users : hrt_5
    resources : database - test_1, table - table_1, column -
c0
    access - all access denied
  data setup :-
    database - test_1
    table - table_1(c0 int, c1 int)
```

The user is able to insert data into the table.

None.

Known Issues in Schema Registry

Learn about the known issues in Schema Registry, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

CDPD-40380: Authorization checking issue when Kerberos is disabled

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the `ranger.plugin.schema-registry.service.name` field.
3. Replace `GENERATED_RANGER_SERVICE_NAME` with the actual name of the service.
4. Restart the Schema Registry service.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

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

Schema Registry uses a schema fingerprinting mechanism to differentiate between schemas. However, fingerprinting does not take into consideration the schema attributes of the field type. As a result, if you have two schemas where the only difference is that one has type attributes defined and the other does not, they will be considered identical by Schema Registry. For example, the following schemas are considered identical:

```
#Schema V1
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [{"name": "local_timestamp_micros_long", "type": "long"}]}

#Schema V2
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [{"name": "local_timestamp_micros_long", "type": {"type": "long", "logicalType": "local-timestamp-micros"}}]}
```

Notice that the only difference is that in the second schema, the `local_timestamp_micros_long` field has a logical type specified. In cases like this, the new version of the schema is not created, the initial version is used. This is true even if the data that is being produced has a new schema version. The ID of the first schema version is used and is put in the serialized record. The new schema version is not created.

This issue is common when using change data capture (CDC) connectors like the Debezium connectors. This is because CDC connectors create schemas with the logical type decimal based on the column type in the database schema. For example:

```
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [{"name": "database_column", "type": {"type": "bytes", "logicalType": "decimal", "precision": 64, "scale": 0}}]}
```

If the database schema changes (for example, the column type), it is possible that only scale changes, which is a schema attribute.

```
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [{"name": "database_column", "type": {"type": "bytes", "logicalType": "decimal", "precision": 64, "scale": 1}}]}
```

In this case, even though scale changed to 1, the first version of the schema is used where scale is 0. As a result, the data is consumed with the wrong scale.

Avoid using logical types or other attributes. Alternatively, ensure that there are no changes in the logical types or other attributes between schema versions.

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

Schema Registry does not start if the address specified in the Load Balancer Address Ranger property does not end with a trailing slash (/).

Set the value of the RANGER_REST_URL Schema Registry environment variable to an address that includes a trailing slash.

1. In Cloudera Manager, select the Schema Registry service.
2. Go to Configuration.
3. Find the Schema Registry Server Environment Advanced Configuration Snippet (Safety Valve) property and add the following:

```
Key: RANGER_REST_URL
Value: [***RANGER REST API URL***]
```

Replace [***RANGER REST API URL***] with an address that can be used by Schema Registry to access Ranger. Ensure that the address ends with a trailing slash. For example: `https://ranger-1.cloudera.com:6182/`

4. Restart the Schema Registry service.

OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

1. Access the Ranger Console (Ranger Admin web UI).
2. Click the cm_atlas resource-based service.
3. Add the schemaregistry user to the all - * policies.
4. Click Manage Service Edit Service .
5. Add the schemaregistry user to the default.policy.users property.

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

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Known Issues in Spark

Learn about the known issues in Spark the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14**CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files**

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the RuntimeWarning: Failed to add file message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

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

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

The following services support DDL statements:

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

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

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

CDPD-67517: Spark3 tests are failing if /tmp is mounted as NOEXEC.

Map the tmpdir to a writable path in spark3-conf/spark-defaults.conf using the following steps:

1. Go to the Cloudera Manager.
2. Go to the Spark service.
3. Click the Configuration tab.
4. Select Scope Gateway .
5. Select Category Advanced .
6. Locate the Spark Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-defaults.conf_client_config_safety_valve property.
7. Map the tmpdir to a writable path:

```
spark.driver.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
spark.executor.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
```

8. Enter a Reason for change, and then click Save Changes to commit the changes.
9. Deploy the client configuration.

CDPD-23817: In the upgraded Cluster, the permission of /tmp/spark is restricted due to the HDP configuration hive.exec.scratchdir=/tmp/spark.

If you are using the /tmp/spark directory in the CDP cluster, you must provide the required additional Policies/ACL permissions.

CDPD-22670 and CDPD-23103: There are two configurations in Spark, "Atlas dependency" and "spark_lineage_enabled", which are conflicted. The issue is when Atlas dependency is turned off but spark_lineage_enabled is turned on.

Run Spark application, Spark will log some error message and cannot continue. That can be restored by correcting the configurations and restarting Spark component with distributing client configurations.

CDPD-23007: Mismatch in the Spark Default DB Location. In HDP 3.1.5, hive_db entities have one attribute - 'location' which is configured to the '/managed' path. In fresh install of CDP 7.1.7, hive_db entities now have 2 attributes 'location' configured to '/external' path and 'managedLocation' configured to '/managed' path. In the AM2CM migration (HDP 3.1.5 -> CDP 7.1.7), the 'location' attribute from hive_db entities in HDP 3.1.5 comes unaltered to CDP 7.1.7 and hence maps to '/managed' path.

This issue arises only if you are upgrading from HDP 3.1.5 to CDP 7.1.7. If you are performing a fresh install of CDP 7.1.7, you can ignore this issue.

None

CDPD-217: The Apache Spark connector is not supported

The old *Apache Spark - Apache HBase Connector* (shc) is not supported in CDP releases.

Use the new HBase-Spark connector shipped in CDP release.

CDPD-3038: Launching pyspark displays several HiveConf warning messages

When pyspark starts, several Hive configuration warning messages are displayed, similar to the following:

```
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of name hive.meta
store.runworker.in does not exist
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of
name hive.masking.algo does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of
name hive.metastore.runworker.in does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of na
me hive.masking.algo does not exist
```

These errors can be safely ignored.

Known Issues in Sqoop

Learn about the known issues in Sqoop, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version

If you have upgraded to CDP PvC Base 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to CDP PvC Base 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see [Troubleshooting Apache Sqoop issues](#).

CDPD-44431: Using direct mode causes problems

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the `--direct` option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the `sqoop.enable.deprecated.direct` property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through `-Dsqoop.enable.deprecated.direct=true`.

CDPD-3089: Avro, S3, and HCat do not work together properly

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

Parquet columns inadvertently renamed

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.

PARQUET-99: Importing Parquet files might cause out-of-memory (OOM) errors

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.

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 in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager

Cloudera Manager does not display a Log Files menu for SMM UI role (and SMM UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by SMM UI is not supported by Cloudera Manager.

View the SMM UI logs on the host.

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in SMM safety valve in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Add the following value for bootstrap servers>Save Changes > Restart SMM.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

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

KNOX does not send neither the Authentication header nor the `hadoop.auth` cookie. Because of this, the SMM UI returns an HTTP 401 response and sets the `www-authenticate: Negotiate` header. After this, KNOX still does not send the Authentication header. This results in the SMM UI being inaccessible through Knox.

You can access the SMM UI directly in Cloudera Manager at `Clusters SMM Streams Messaging Manager Web UI` and log in using the Kerberos username and password.

CDPD-82560: Streams Messaging Manager known unprotected endpoint

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

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

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 in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

CDPD-11079: Blacklisted topics appear in the list of replicated topics

If a topic was originally replicated but was later excluded for replication, 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, this topic will be included in the response. Additionally, the excluded topic will also be visible in the SMM UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

CDPD-30275: SRM may automatically re-create deleted topics on target clusters

If auto.create.topics.enable is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with srm-control. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1]
```

2. Wait until SRM is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

Limitations

SRM cannot replicate Ranger authorization policies to or from Kafka clusters

Due to a limitation in the Kafka-Ranger plugin, SRM cannot replicate Ranger policies to or from clusters that are configured to use Ranger for authorization. If you are using SRM to replicate data to or from a cluster that uses Ranger, disable authorization policy synchronization in SRM. This can be achieved by clearing the Sync Topic Acls Enabled (sync.topic.acls.enabled) checkbox.

SRM cannot ensure the exactly-once semantics of transactional source topics

SRM data replication uses at-least-once guarantees, and as a result cannot ensure the exactly-once semantics (EOS) of transactional topics in the backup/target cluster.



Note: Even though EOS is not guaranteed, you can still replicate the data of a transactional source, but you must set isolation.level to read_committed for SRM's internal consumers. This can be done by adding `[***SOURCE CLUSTER ALIAS***]->[***TARGET CLUSTER ALIAS***].consumer.isolation.level=read_committed` to the Streams Replication Manager's Replication Configs SRM service property in Cloudera Manager. The isolation.level property can be set on a global connector or replication level. For example:

```
#Global connector level
connectors.consumer.isolation.level
=read_committed
#Replication level
uswest->useast.consumer.isolation.l
evel=read_committed
```

SRM checkpointing is not supported for transactional source topics

SRM does not correctly translate checkpoints (committed consumer group offsets) for transactional topics. Checkpointing assumes that the offset mapping function is always increasing, but with transactional source topics this is violated. Transactional topics have control messages in them, which take up an offset in the log, but they are never returned on the consumer API. This causes the mappings to decrease, causing issues in the checkpointing feature. As a result of this limitation, failover operations for transactional topics is not possible.

Known Issues in YARN 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 in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

None.

COMPX-4644: Queue capacity rounding problem when configuration is initially set via YARN

When setting the capacity scheduler configuration through the YARN/Cloudera Manager configuration, there may be capacity values that use multiple decimal places. This results in rounding/floating point precision discrepancies in the UI when trying to validate that all sibling capacities equal 100%. The UI looks like all the numbers add up to 100, but the validation still displays an error and does not allow to save the capacities. It is also observed that the capacity is being calculated as, for example, 99.9999999991 in the backend.

- Create queues within the UI, or

- Ensure that capacities configured through the Capacity Scheduler safety valve do not have more than one decimal place.

20202 Database migration after enabling opt-in migration

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to CDP 7.1.9, forcing migration from H2 to PostgreSQL database.
- Upgrade to CDP 7.1.9 CHF2, moving back to H2 database.
- Upgrade to CDP 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version CDP 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

CDPD-56559: MapReduce jobs can intermittently fail during a rolling upgrade.

During a rolling upgrade between CDP versions 7.1.8 and 7.1.9, MapReduce jobs may fail with message, `RuntimeException: native snappy library not available`. Although the native Snappy compression library is not loaded, a checkmark displays to indicate that the Snappy compression library is loading for NodeManagers that are pending upgrades. This causes the MapReduce jobs that are associated with the NodeManagers to fail. After the upgrade, the jobs work as expected. This issue only impacts rolling upgrades from before CDP 7.1.9 to a higher version.

None.

COMPX-12021 Queue Manager configurations on Scheduler Configuration page are not working

When setting the following properties on the YARN Queue Manager UI, the properties are set in the `capacity-scheduler.xml` file which does not have any effect on YARN. The properties need to be set in the `yarn-site.xml` file, which does not happen when you set them through YARN Queue Manager.

- "Maximum Application Priority" – `"yarn.cluster.max-application-priority"`
- "Enable Monitoring Policies" – `"yarn.resourcemanager.scheduler.monitor.enable"`
- "Monitoring Policies" – `"yarn.resourcemanager.scheduler.monitor.policies"`



Note: You can also set this property on the YARN configuration page in Cloudera Manager as "Capacity Scheduler Preemption".

- "Preemption: Observe Only" – `"yarn.resourcemanager.monitor.capacity.preemption.observe_only"`
- "Preemption: Monitoring Interval (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.monitoring_interval"`
- "Preemption: Maximum Wait Before Kill (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.max_wait_before_kill"`
- "Preemption: Total Resources Per Round" – `"yarn.resourcemanager.monitor.capacity.preemption.total_preemption_per_round"`
- "Preemption: Over Capacity Tolerance" – `"yarn.resourcemanager.monitor.capacity.preemption.max_ignored_over_capacity"`
- "Preemption: Maximum Termination Factor" – `"yarn.resourcemanager.monitor.capacity.preemption.natural_termination_factor"`
- "Enable Intra Queue Preemption" – `"yarn.resourcemanager.monitor.capacity.preemption.intra-queue-preemption.enabled"`

1. In Cloudera Manager, select the YARN service.
2. Click the Configuration tab.
3. Search for `yarn-site.xml`.
4. Under YARN Service Advanced Configuration Snippet (Safety Valve) for `yarn-site.xml`, add the corresponding parameter and value you need.

5. Click Save Changes.
6. Restart the YARN services.

COMPX-13177: QueueManager webapp requests fail with 'HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA'

Products:

- Cloudera Manager for CDP Private
- Cloud Base Cloudera Manager for CDP Public Cloud

Context:

- Centos 7.8 and Redhat 7.8 operating systems, when FIPS support is enabled.

Problem:

- When attempting to display the Yarn Queue Manager interface, Cloudera Manager displays an error: "HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA".
1. Edit the file [/etc/default/cloudera-scm-server]
 2. Around line 28, modify the line that starts with

```
#export CMF_OVERRIDE_TLS_CIPHERS=...
```

3. Remove the comment mark #.
4. Remove all ciphers with "3DES" in the name.
5. Save the file.
6. Restart the Cloudera Manager Server service.

Third-party applications do not launch if MapReduce framework path is not included in the client configuration

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the `mapreduce.application.framework.path` property is set to the appropriate value, but third-party applications with their own configurations does not launch.

Set the `mapreduce.application.framework.path` property to the appropriate configuration for third-party applications.

JobHistory URL mismatch after server relocation

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated the `mapred-site.xml` file that references the correct JobHistory Server.

YARN cannot start if Kerberos principal name is changed

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN does not start. In such cases, the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the `znode` and restart the YARN service.
- Use the `reset ZK ACLs` command. This also sets the `znodes` below `/rmstore/ZKRMStateRoot` to `world:anyone:cdw` which is less secure.

Queue Manager does not open on using a custom user with a default Kerberos principal

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager process_username property matches the YARN process_username property.

Known Issues in Zeppelin

Learn about the known issues in Zeppelin the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

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.

BUG-125263: Zeppelin service move fails on clusters upgraded from HDP3.1.5

Resolve the circular symlink issue on the Zeppelin node by linking the conf directory to a new directory under /etc/zeppelin:

- # mkdir -p /etc/zeppelin/<version>/0
- # rm /usr/hdp/<version>/zeppelin/conf
- # ln -s /etc/zeppelin/<version>/0 /usr/hdp/<version>/zeppelin/conf

Where version is the HDP version. For example, 7.1.x-yyy. Restart the Zeppelin server in Ambari.

CDPD-3090: Due to a configuration typo, functionality involving notebook repositories does not work

Due to a missing closing brace, access to the notebook repositories API is blocked by default.

From the CDP Management Console, go to Cloudera Manager for the cluster running Zeppelin. On the Zeppelin configuration page (Zeppelin serviceConfiguration), enter shiro urls in the Search field, and then add the missing closing brace to the notebook-repositories URL, as follows:

```
/api/notebook-repositories/** = authc, roles[{{zeppelin_admin_group}}]
```

Click Save Changes, and restart the Zeppelin service.

CDPD-2406: Logout button does not work

Clicking the Logout button in the Zeppelin UI logs you out, but then immediately logs you back in using SSO.

Close the browser.

OPSAPS-59802: Zeppelin and Livy roles should be co-located on the same host.

When installing or upgrading to CDP Private Cloud Base, you must co-locate all Zeppelin and Livy roles on the same cluster host due to an issue with certificate generation.

Known Issues in ZooKeeper

Learn about the known issues in ZooKeeper the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 14.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 14

There are no new known issues identified in the previous releases.

Cumulative hotfix 13

Know more about the cumulative hotfix 13 for CDP 7.1.9 SP1. This cumulative hotfix was released on December 15, 2025.

Cloudera Runtime 7.1.9.1066-8 (Cumulative Hotfix 13) download URL:

Parcel Repository Location

https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1066/parcels/

What's new in Cloudera Runtime 7.1.9 SP1 CHF 13

Understand the functionalities and improvements to features of components in Cloudera Runtime 7.1.9 SP1 CHF 13.

Atlas

The label for classification count is changed in Statistics

In the **Statistics** page, the **Count** label is changed to **Associated Entities** to avoid confusion. This number shows the number of entities marked with that classification.

Statistics

Classifications (24)

Name	Associated Entities (24)
Aggregated	1
Airline Alliances	1
Archive	1
Business Domain	2
Business Performance	2
Enriched	1
Financial_Settlement	2
Fleet Management	2
Flight Operations	1
GDPR_PNR	4
Master	1
Operations	1
PII_Combined	1
Raw	1
Revenue Management	2
Staged	1

Close

Fixed issues in 7.1.9 SP1 CHF 13

Know more about the list of fixes that are shipped for CDP Private Cloud Base version 7.1.9 SP1 CHF 13.



Note: Contact Cloudera Support for questions related to any specific hotfixes. The version and build identifier for this CDH Parcel is 7.1.9-1.cdh7.1.9.p1066.73837803.

CDPD-64874: Intermittent failure in TestOzoneRpcClientAbstract.testListSnapshot

The `ListSnapshot` API uses the `org.apache.hadoop.ozone.om.ListIterator.MinHeapIterator` static class, which internally uses the `CacheIterator` and `DBIterator` objects. The `DBIterator` object checks if the rocks DB key is present in the cache in the `org.apache.hadoop.ozone.om.ListIterator.DbTableIter#getNextKey` method. Previously, this check accessed the cache from the table cache, which could be intermittently flushed, causing the addition of duplicate entries in the `org.apache.hadoop.ozone.om.ListIterator.MinHeapIterator` static class. This issue is now fixed and you must use the preloaded keys in the `org.apache.hadoop.ozone.om.ListIterator.CacheIter#cacheKeyMap` method within the `org.apache.hadoop.ozone.om.ListIterator.CacheIter` static class.

Apache Jira: [HDDS-9967](#)

CDPD-69064: DataNode volume DU reserved percentage should have a non-zero default value

Previously, the `hdds.datanode.dir.du.reserved.percent` `DataNode` configuration default value was zero. This issue is fixed now and the default value is 0.0001, which means that the reserved percentage is 0.01%. As a result, for each `DataNode` volume being used by Ozone, the available volume capacity for Ozone is calculated by `real capacity volume - (real capacity of the volume * 0.0001)`, ensuring that 0.01% of the volume is preserved for non-Ozone use.

Apache Jira: [HDDS-10720](#)

CDPD-90382: FSO to FSO replication failing for snapshot diff issue

Previously, FSO to FSO bucket replication using Replication Manager failed if the ozone sh snapshotdiff report between the current and previous snapshot contained more than 1000 entries. This was because Replication Manager could not parse the additional newline at the end of each page in the paginated response. This issue is now fixed by removing the additional newline at the end of every page in snapshotdiff output.

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

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

CDPD-87831: Avoid scheduling replications on full DataNodes by tracking pending operation size in SCM

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

Apache Jira: [HDDS-13437](#)

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

Previously, if the leader node in pipeline did not have capacity to create a new container, it might have returned a container creation failure. If the follower node did not have the capacity to create new container, it might have failed and kept trying for another follower node to be successful. This could cause the disk to fill up with parallel write blocks through a state machine, and slowed down the write capability and failure response. This issue is now fixed by checking whether a `DataNode` has enough space for a new container before allocating a new container to that `DataNode`. This improves write performance and reduces container creation failure in scenarios when `DataNodes` have less than 5GB disk space remaining.

Apache Jira: [HDDS-12468](#)

CDPD-80031: Missing latency metrics in Ozone Manager for deletion services

Previously, no metrics were available in Ozone Manager to measure how long each iteration of the deletion services were taking. This issue is now fixed and latency metrics for key deletion services of Ozone Manager are available now.

Apache Jira: [HDDS-12442](#)

CDPD-85199: Upgraded jackson-core to 2.15.0

The jackson-core version of the Jackson Data Processor for Phoenix Query Server (PQS) and OMID is upgraded to 2.15.0 due to CVE-2025-49128 and CVE-2025-52999.

CDPD-76594: Impala Iceberg table creation fails in Hue

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

CDPD-78292: Hue defunct gunicorn worker processes

On Ubuntu 22, when using Oracle Database, `rungunicornserver` worker processes might have accumulated due to incomplete process termination and stale database connections. This resulted in a cluttered process table but did not critically impact service functionality. This issue is now resolved.

CDPD-90930: Hue Metadata front end cache not functioning correctly

Previously, the metadata front end cache was not applied as expected, resulting in slower page load times and degraded responsiveness. This issue is now resolved, and metadata is properly cached, improving UI performance and overall user experience.

CDPD-89886: FileUtils.isOwnerOfFileHierarchy does not have the recurse parameter

This patch resolves performance degradation in `add_partition` API calls to the Hive Metastore (HMS). The latency of the `add_partition` calls increases significantly with the number of folders and files in HDFS. You can optimize performance by using the existing `xasecure.hive.ur i.permission.coarse.check` configuration parameter. Setting this value to true prevents recursive permission checks on the parent folder.

CDPD-89263: Enhanced file extension controls for Hue file management

Previously, Hue allowed only the control of allowed and restricted file type extensions for Hue file uploads.

This issue is now fixed by Hue allowing to control both allowed and restricted file type extensions for Hue file uploads and create and rename the file. This improves flexibility and enhances the overall user experience.

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

CDPD-87409: Duplicate compaction processing by multiple HiveServer2 threads

Multiple HiveServer2 Compactor Worker threads could process the same minor compaction simultaneously. This created duplicate delta directories in HDFS and caused subsequent table operations to fail with an `IllegalStateException` error.

This issue is now resolved by implementing a synchronization mechanism. The compaction worker now locks the compaction ID in the metadata store before starting, ensuring that only one thread processes a compaction at a time.

CDPD-91034: MariaDB falls back to MySQL in Hive

Hive downstream had errors in supporting MariaDB.

The issue is now fixed by making MariaDB automatically fall back to MySQL.

CDPD-92487: Impala compatibility with OpenSSL 3.2 on RHEL 9.5

Impala experienced widespread failures, including numerous test failures and cluster hangs, after RHEL/Rocky 9.5 updated its OpenSSL version from 3.0.2 to 3.2.x. This OpenSSL rebase broke two core areas in Impala, affecting both AES-GCM encryption and SSL connection handling.

The issue is now resolved by correcting the OpenSSL call sequence for AES-GCM encryption (IMPALA-13592). The Initialization Vector (IV) length is now set before the IV value, ensuring that encryption works properly with OpenSSL 3.2 and by avoiding the unnecessary flush() call when the TSSLSocket transport is closing after an SSL handshake error (IMPALA-13680). This prevents the indefinite hang and ensures that failed connections are correctly terminated.

Apache Jira: [IMPALA-13592](#),[IMPALA-13680](#)

CDPD-91626: Stable file system operations

Unreadable paths in Java CLASSPATH could lead to crashing during startup without an informative error message. This is now resolved by ignoring unreadable paths in CLASSPATH.

Apache Jira: [IMPALA-13371](#)

CDPD-81958: Improve cookie security in Knox-proxied web UIs with Secure and HttpOnly attributes

Previously, the pac4jCsrfToken cookie in Apache Knox-proxied applications did not contain the Secure and HttpOnly flags. This issue is now fixed by ensuring that the pac4jCsrfToken cookie has both the Secure and HttpOnly flags in Knox-proxied applications, improving the security provided by Knox.

Apache JIRA: [KNOX-3134](#)

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

CVE-2014-0114	commons-beanutils
CVE-2019-10086	commons-beanutils
CVE-2024-52046	mina-core
CVE-2025-59059	Apache Ranger
CVE-2025-59060	Apache Ranger

Behavioral changes in Cloudera Runtime 7.1.9 SP1 CHF 13

Know more about the behavioral changes in cumulative hotfix 13 for 7.1.9 SP1.

Hive

CDPD-92002: New configuration properties added for Hive metatool summary operation

A new configuration property is available to manage the performance and scope of the metatool summary operation, particularly when dealing with the task running indefinitely.

Previous behavior:

The Hive metatool summary operation executed without specific controls for resource usage, time limits.

New behavior:

A new configuration property is available to manage the performance and scope of the metatool summary operation, particularly when dealing with the task running indefinitely.

Property	Default Value	Description
hive.metatool.summary.timeout	20 min	The maximum time the summary operation will wait to complete before being automatically canceled.

Impala

CDPD-89269: Batched Insert events for Impala

Previous behavior:

Impala initiated insert statements using a single API call for each partition. This could become a performance bottleneck when inserting data into tables with a large number of partitions.

New behavior:

This change speeds up computing file checksums during **INSERT** operations, especially for tables with many files and partitions. This is achieved by leveraging a new batch insert API in the Hive Metastore to process events more efficiently.

Apache Jira: [IMPALA-14051](#)

CDPD-92378: Impala no longer allows unauthenticated user impersonation through Knox JDBC gateway

Previous behavior:

When connecting to Impala using the Knox JDBC gateway URL (e.g., jdbc:impala://[knoxhost]:[knoxport]/...), you could specify any Impala user through the `impala.doas.user` property and successfully impersonate them, regardless of whether your Knox user had explicit Impala proxy user permissions. This unintended behavior bypassed security checks and allowed for delegation to users who should not have been accessible.

New behavior:

Knox JDBC Gateway authenticates as itself and delegates to the logged in user, so it has already performed some delegation. If `impala.doas.user` is passed by the client through Knox JDBC Gateway, it will now result in an error as we don't support secondary delegation.

Known issues in 7.1.9 SP1 CHF 13

The following section lists the known issues and limitations, the areas of impact, and workaround for each Runtime component.

Known Issues in Atlas

Learn about the known issues in Atlas the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

CDPD-93307: Hive import logs are truncated during bulk import

When running the `import-hive.sh` script to perform bulk Hive metadata imports, the logging process fails to retain the expected history. The `import-hive.log` file rotates prematurely when it reaches 10 MB and retains only one backup file (`import-hive.log.1`), instead of the expected 100 MB file size and 20 backup files.

Supply a custom `log4j` configuration file with the correct structure by using the `ATLAS_OPTS` environment variable before running the import script:

```
<?xml version="1.0" encoding="UTF-8" ?>
<!DOCTYPE log4j:configuration SYSTEM "log4j.dtd">
<log4j:configuration xmlns:log4j="http://jakarta.apache.org/log4j/">
  <appender name="console" class="org.apache.log4j.ConsoleAppender">
    <param name="Target" value="System.out"/>
    <layout class="org.apache.log4j.PatternLayout">
      <param name="ConversionPattern" value="%d %-5p - [%t: %x] ~ %m (%C{1}:%L)%n"/>
    </layout>
  </appender>
  <appender name="FILE" class="org.apache.log4j.RollingFileAppender">
    <param name="File" value="${atlas.log.dir}/${atlas.log.file}"/>
    <param name="Append" value="true"/>
```

```

        <param name="maxFileSize" value="100MB"/>
        <param name="maxBackupIndex" value="20"/>
        <layout class="org.apache.log4j.PatternLayout">
            <param name="ConversionPattern" value="%d %-5p - [%t:
%x] ~ %m (%C{1}:%L)%n"/>
        </layout>
    </appender>
    <logger name="org.apache.atlas" additivity="false">
        <level value="info"/>
        <appender-ref ref="FILE"/>
    </logger>
    <logger name="org.apache.kafka.common.config.AbstractConfig"
additivity="false">
        <level value="error"/>
        <appender-ref ref="FILE"/>
    </logger>
    <root>
        <priority value="warn"/>
        <appender-ref ref="FILE"/>
    </root>
</log4j:configuration>

```

Export the ATLAS_OPTS variable to point to your custom file before running the import script:

```

export ATLAS_OPTS="-Dlog4j.configuration=file:/path/to/custom-hi
ve-import-log4j.xml"
./import-hive.sh

```

CDPD-87670: [Atlas UI] Apache Atlas Glossary becomes unresponsive when the page size is set to 50

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

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-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-6565: Issue with ddlQueries in Atlas for Hive/Impala tables created in default/non-default db using DAS/HUE respectively

ddlQueries are not created for Impala origin table. ddlQueries created for Impala/Hive CTAS tables have difference in names as compared to same operation for Hive table. The Impala table has the query in the ddlQueries entity name.

None

CDPD-55301: ddlQueries and ALTERNATIVE_* lineage missing for Spark tables created through spark3-shell

The ddlQueries and ALTERNATIVE_* lineage missing for Spark tables created using spark3-shell.

None

CDPD-56085: 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

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*`.

None

CDPD-58554: Discard audits of specific classification, label, and business metadata

Support to control audits for specific classification, label, and business metadata is not present in the Custom Audit Filters feature.

None

CDPD-40346: The `ddlQueries` and `ALTERTABLE_ADDCOLS` lineage missing for Impala tables.

The `ALTERTABLE_ADDCOLS` lineage has some issue when an Impala table is altered and the corresponding lineage is not created.

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

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

None

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

Adding Japanese labels or user defined properties results in the error message: "Invalid label: ###, label should contain alphanumeric characters, _ or -"

None

CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties

Entities, such as `hdfs_path`, are not returned by the search when their indexable properties are searched using partial search terms made of Japanese or Chinese characters. Only exact matches return results when searching the indexed properties made of Japanese or Chinese characters.

None

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

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

None.

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

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

None.

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

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

None.

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

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

None.

CDPD-43772: Performance issues with Atlas service

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

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

OPSAPS-67783: During rolling upgrade one among two Atlas server failed to start but Cloudera Manager considered as success

Cloudera Manager marks the Execute command Start on service Atlas-1 as a success even when Atlas service had failed to start successfully. In such cases, Atlas logs give the exact reason for Atlas start-up failure.

None

CDPD-19358: "IsIndexable" and "isOptional" value of a typedef's attribute is modified post migration.

In HDP-265, falcon_feed_creation has "stored-in" attribute which has IsIndexable value set to True.

Post migration, "stored-in" is moved to relationshipAttributeDefs and has IsIndexable value set to False.

Similarly, in HDP-265 the table attribute of hive_storage_desc isOptional is set to true. Post migration, isOptional set to False.

None.

CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command

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

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

In converting content from Navigator to Atlas, the create time for Hive tables is not moved to Atlas.

None

CDPD-10574: Suggestion order doesn't match search weights

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

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@58112
bc4 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) does not 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 the metadata collection from HMS. For example, if an ALTER TABLE operation in Spark that is changing a table name and is reporting 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 "@" does not 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, and use 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 does 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 logout 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 CDP.

None.

CDPD-1664: Guest users are redirected incorrectly

Authenticated users logging in to Atlas are redirected to the CDP 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: The 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, to add metadata tags to relationships that you wanted to make sure were unique in the system, you can design the relationship attribute with the property IsUnique equal true. However, in this release, the IsUnique attribute is not enforced.

None.

DOCS-13759: Tag Propagation stops after a certain depth while the lineage is being extended

When a tag is added to an entity at timestamp T1, the entities along the lineage to which the tag must be propagated is calculated at T1. Before tag propagation completes, if the lineage is extended, tag does not propagate to the entities in the extended lineage.

CDPD-41142: When a Kafka console consumer group is run, more than one update audits are seen

After running the console consumer with a consumer group, verify the consumer group entity created, along with the metrics and notifications for the consumer group and topic. The expected result can be: one ENTITY_CREATE audit and one ENTITY_UPDATE audit. But more than one ENTITY_UPDATE audits are seen.

CDPD-40165: Two audits are created for SPARK CTAS table

When following Spark queries are fired:

```
spark.sql("create table table1(id int)")
spark.sql("create table table2 as select * from table1")
```

HMS sends "ENTITY_CREATE" and "ENTITY_FULL_UPDATE_V2".

The extra ENTITY_FULL_UPDATE_V2 message received from HMS is sent as a part of ALTER_TABLE_ADDCOLS event from the HMS Hook side. This behaviour is observed only when the queries are run from Spark SQL and not when run the same queries are run from Beeline.

CDPD-39197: Debug metrics returns empty data

When debug metrics is enabled and some operations are performed, the response is empty

CDPD-36495: Updating legacyAttribute from False to True resets the initially created relationshipAttributes values

Creating types, entities, and to start, you must set the relationship with is_legacy_attribute value as False.

Later, update the value relationshipDef is_legacy_attribute to True.

For the entities that were created before updating the is_legacy_attribute to True, the relationship Attributes value is reset.

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-29409: Hive import: Suggestion suggests entity which is deleted.

Suggestions suggests tables of a database, which is a deleted entity.

CDPD-25152: Tag propagation through deferred actions consumes additional time as compared to default flow

The additional time might be due to the small overhead added to create / update task vertex and which is run in the background. This also depends on number of tasks queued to be executed.

CDPD-42954: Zeppelin notebook fails after enabling Atlas-HDFS hook

The Zeppelin notebooks are failing with errors after enabling Atlas-HDFS hook in the CDP cluster.

When the below properties are set for atlas-client.properties in Cloudera Manager:

- atlas.jaas.KafkaClient.option.keyTab
- atlas.jaas.KafkaClient.option.principal

Along with adding the properties in /etc/atlas/conf/atlas-application.properties, Cloudera Manager also adds these properties to atlas-application.properties for other services (like Spark).

Adding these properties interferes with the normal flow of the services (like Spark)

To enable HDFS lineage feature, instead of setting these properties through Cloudera Manager, users can manually add the properties directly in /etc/atlas/conf/atlas-application.properties

CDPD-10576: Deleted Business Metadata attributes appear in Search Suggestions

Atlas search suggestions continue to show Business Metadata attributes even if the attributes have been deleted.

None.

CDPD-39427: [HDFS Lineage]: When the input is a directory in case of put/copyfromLocal/cp/mv commands, lineage is not created even though the script succeeds.

When Source is a directory and target is a directory which is already present in Atlas, the command succeeds and inserts the data in the desired location, but lineage is not created.

DOCS-13760: System Attributes search, __classificationNames: Search with parent tag does not return entities associated to its children tags

System attribute search with __classificationNames = parent_tag returns entities associated to parent_tag only and not entities associated to its children tag.

Workaround: Instead of using system attribute, employ the basic search attribute "classification" which lists entities associated with inherited classifications.

CDPD-35818: Basic search with tag filter provides approximateCount as -1 when there is no match and is 0 otherwise

When the following search operations are performed:

- Faceted search with both tag filter and entity filter

The observed approximateCount is -1.

Here both entity filter and tag filter are present and when there is no match the response received is -1.

- Faceted search with only entity filter

Performing a basic search provides an approximate value of 0 when there is no match.

- Faceted search with only tag filter

Whenever there is tag filter in the query and there is no entity match, the approximateCount is -1 and if the tag filter is not available, the response approximateCount is 0

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet `atlas.authentication.method.file` is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export

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

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](#)

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

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

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

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

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

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

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the `PermanentLockingException`.

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

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-88230: Bulk import operations currently do not construct lineage as expected

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

Bulk import operations do not create audits for the imported entities.

Known Issues in Cruise Control

Learn about the known issues in Cruise Control, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

Rebalancing with Cruise Control does not work due to the metric reporter failing to report the CPU usage metric

On the Kafka broker, the Cruise control metric reporter plugin may fail to report the CPU usage metric.

If the CPU usage metric is not reported, the `numValidWindows` in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.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 the [Moving a Host Between Clusters](#).



Note: Cluster nodes affected by this issue are not displayed as unhealthy.

Known Issues in Hive

Learn about the known issues in Hive, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

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

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

The following services support DDL statements:

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

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

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

CDPD-60770: Beeline Authentication Issue with Special Characters in Passwords

When LDAP is enabled, users cannot authenticate with Beeline if the password contains a special character. For example, the following string fails:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>;user=user@XXX;password='R3G#xpXyoy1M0Jb1'
```

Use the -p parameter to execute the Beeline command:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>; -n user@XXX -p 'R3G#xpXyoy1M0Jb1'
```

CDPD-43769: Inconsistent behavior with CHAR comparisons and string literals

In Hive, 'char' data to 'string' data comparison is considered as equal, but 'char' data to 'varchar' data comparison is considered as not equal if the data has trailing spaces. For example,

```
create table test(colchar char(10), colvarchar varchar(20), colstring string);
insert into test values ('a ', 'a ', 'a ');

select count(*) from test where colchar = colstring;
Output returns 1 because char and string data is considered as equal

select count(*) from test where colchar = colvarchar;
Output returns 0 because char and varchar data is considered as not equal
```

It is recommended that you use 'varchar or string' data type instead of 'char' data type for any character or string representation.

CDPD-42726: Wrong results for agg queries from stats after running multi-insert query on managed table

Multiple inserts at the same time into the same partition results in invalid stats. For example:

```
create table source(p int, key int, value string);
insert into source(p, key, value) values (101,42,'string42');

create table stats_part(key int,value string) partitioned by (p
int);
from source
insert into stats_part select key, value, p
insert into stats_part select key, value, p;

select count(*) from stats_part;
-- In this scenario, StatsOptimizer helps serving this query be
cause the result should be rowNum of the partition p=101.
-- The result is 1, however, the result should be 2.
```

Rewrite the multi-insert statement by merging the branches inserting into the same partition to one branch.

CDPD-57574: Query execution fails due to NullPointerException in DagUtils.setupQuickStart

Hive queries may fail at runtime with a NullPointerException while executing the Tez graph at DagUtils.setupQuickStart method. A part of the stack trace from the error is shown below:

```
ERROR : Failed to execute tez graph.
java.lang.NullPointerException: null at org.apache.hadoop.hive.ql.exec.tez.DagUtils.setupQuickStart(DagUtils.java:1724)
```

The error usually occurs from an invalid plan that has cycles created by semi-join or map-join edges.

Disable some of the optimizations in order to remove the cycles from the plan. It is sufficient to set one of the following properties:

- set hive.tez.dynamic.partition.pruning=false;
- set hive.tez.dynamic.semijoin.reduction=false;
- set hive.auto.convert.join=false;

To avoid any undesired performance regressions, disable the optimization(s) only for specific queries and not globally (through Cloudera Manager) for the whole workload or cluster.

CDPD-45370: Cartesian join fails with an empty table

When the configuration hive.tez.cartesian-product.enabled is set to true, query fails if one of the tables involved in a Cartesian join has zero rows. This results in an IllegalArgumentException with the message Number of items is 0. Should be positive.

To address this issue, set hive.tez.cartesian-product.enabled to false before running the query.

Apache Jira: [\[HIVE-26582\]](#)

CDPD-43107: SemanticException for INSERT INTO statement when Hive Cost-based Optimizer (CBO) is disabled

If you are running the INSERT INTO query and hive.cbo.enable is set to "false", the query fails with a SemanticException. For example,

```
set hive.cbo.enable=false;

CREATE TABLE mytable (
```

```
id INT,
str STRING
);

INSERT INTO mytable (id, str) VALUES (1, 'a');
```

Output:
 org.apache.hadoop.hive.ql.parse.SemanticException: 0:0 Expected 2 columns for insclause-0/default@mytable; select produces 1 columns. Error encountered near token 'a'

This issue occurs only when columns are specified for the target table in the query. INSERT INTO mytable values (1, 'a'); does not result in an exception.

Enable CBO and run the query.

CDPD-43957: HiveServer shuts down during replication due to high resource usage

During Hive replication (both bootstrap and incremental), you may notice that the HiveServer (HS2) shuts down periodically with the following error:

```
java.sql.SQLException: org.apache.hive.jdbc.ZooKeeperHiveClientException: Unable to read HiveServer2 configs from ZooKeeper
    at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.java:265)
    at org.apache.hive.jdbc.HiveDriver.connect(HiveDriver.java:107)
    at java.sql.DriverManager.getConnection(DriverManager.java:664)
    at java.sql.DriverManager.getConnection(DriverManager.java:247)
    at com.cloudera.enterprise.hive3qt.Hive3QueryTool$HiveOperation.execute(Hive3QueryTool.java:682)
    at com.cloudera.enterprise.hive3qt.Hive3QueryTool.main(Hive3QueryTool.java:935)
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Unable to read HiveServer2 configs from ZooKeeper
    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:177)
    at org.apache.hive.jdbc.Utills.configureConnParamsFromZooKeeper(Utills.java:580)
    at org.apache.hive.jdbc.Utills.parseURL(Utills.java:391)
    at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.java:263)
    ... 5 more
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Tried all existing HiveServer2 uris from ZooKeeper.
    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.getServerHosts(ZooKeeperHiveClientHelper.java:132)
    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:172)
    ... 8 more
```

This issue occurs when you replicate at scale with an unbalanced cluster setup that has all the roles running on the Cloudera Manager host. As a result, Cloudera Manager ends up in a bottleneck situation because HS2 crashes, preventing further replication. The logs (/var/log/messages) indicate that there were a large number of Java processes running on the host, which exhausted the host's memory and triggered an Out Of Memory Killer process to stop HS2.

You can restart HS2 and reinitiate replication. However, the replication process may take a longer time to complete.

- Follow these recommendations when setting up your source and target clusters:
 - Ensure that no HS2, Hive metastore (HMS), DataNode, or NameNode roles are running on the Cloudera Manager host.
 - Ensure that you have multiple instances of HS2 and HMS roles on different nodes.
 - It is recommended that you have the HS2 and HMS roles on different nodes than the DataNodes or NameNodes.

These practices increase the possibility of Hive replication at scale completing successfully.

- Limit the scale of Hive replication by temporarily disabling replication policies.

CDPD-40730: Parquet change can cause incompatibility

Parquet files written by the parquet-mr library in CDP 7.1.9, where the schema contains a timestamp with no UTC conversion will not be compatible with older versions of Parquet readers. The effect is that the older versions will still consider these timestamps as they would require UTC conversions and will thus end up with a wrong result. You can encounter this problem only when you write Parquet-based tables using Hive, and tables have the non-default configuration `hive.parquet.write.int64.timestamp=true`.

None.

CDPD-41274: HWC + Oozie issue: Could not open client transport with JDBC Uri

Currently only Spark cluster mode is supported in the Oozie Spark Action with Hive Warehouse Connector (HWC).

Use Spark action in cluster mode.

```
<spark xmlns="uri:oozie:spark-action:1.0">
  ...
  <mode>cluster</mode>
  ...
</spark>
```

CDPD-26556 After an upgrade, querying a CTAS table under certain conditions might throw an exception

If you upgrade your Hive cluster from CDH 6 to CDP 7, create a CTAS table in the CDP cluster from a table you upgraded from CDH, you might see the following exception when you query the new table:

```
class org.apache.hadoop.io.IntWritable cannot be cast to class o
rg.apache.hadoop.hive.serde2.objectinspector.StandardUnionObject
Inspector$StandardUnion
```

This issue involves CDH-based tables having columns of complex types ARRAY, MAP, and STRUCT.

CDPD-23506: OutOfMemoryError in LLAP

Long running spark-shell applications can leave sessions in interactive Hiveserver2 until the Spark application finishes (user exists from spark-shell), causing memory pressure in case of a high number of queries in the same shell (1000+).

You must close spark-shell so that sessions are closed. Add the owner of the database or the tables as a user with read or read/write access to the tables directly.

CDPD-23041: DROP TABLE on a table having an index does not work

If you migrate a Hive table to CDP having an index, DROP TABLE does not drop the table. Hive no longer supports indexes ([HIVE-18448](#)). A foreign key constraint on the indexed table prevents dropping the table. Attempting to drop such a table results in the following error:

```
java.sql.BatchUpdateException: Cannot delete or update a parent
row: a foreign key constraint fails ("hive"."IDXS", CONSTRAINT "
```

```
IDXS_FK1" FOREIGN KEY ("ORIG_TBL_ID") REFERENCES "TBL" ("TBL_ID"
))
```

There are two workarounds:

- Drop the foreign key "IDXS_FK1" on the "IDXS" table within the metastore. You can also manually drop indexes, but do not cascade any drops because the IDXS table includes references to "TBL".
- Launch an older version of Hive, such as Hive 2.3 that includes IDXS in the DDL, and then drop the indexes as described in [Language Manual Indexing](#).

Apache Issue: [Hive-24815](#)

CDPD-17766: Queries fail when using spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive.

Do not specify spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive remotely.

Workaround: specify principal=hive.server2.authentication.kerberos.principal as shown in the following syntax:

```
jdbc:hive://<host>:<port>/<dbName>;principal=hive.server2.authentication.kerberos.principal;<otherSessionConfs>?<hiveConfs>#<hiveVars>
```

CDPD-13636: Hive job fails with OutOfMemory exception in the Azure DE cluster

Set the parameter hive.optimize.sort.dynamic.partition.threshold=0. Add this parameter in Cloudera Manager (Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml)

CDPD-10848: HiveServer Web UI displays incorrect data

If you enabled auto-TLS for TLS encryption, the HiveServer2 Web UI does not display the correct data in the following tables: Active Sessions, Open Queries, Last Max n Closed Queries

Known Issues in HDFS

Learn about the known issues in HDFS, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

There are no known issues identified in the releases before Cloudera Runtime 7.1.9 SP1 CHF 13.

Known Issues in Hue

Learn about the known issues in Hue, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

CDPD-94739: Hue audit logging requires debug mode

In Hue, audit logging is only active when debug logging is enabled by setting DJANGO_DEBUG_MODE = true. By default, audit logging is disabled.

Update to the upcoming release to restore audit logging functionality.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

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

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark

queries in a batch mode launches a shell script Oozie job from Hue and this is not supported on a secure cluster.

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

OPSAPS-69659: Hue service fails on restart with "Unable to find psycopg2 2.5.4" error

Hue service fails to restart and you see the following error: Unable to find psycopg2 2.5.4. This could be because you have installed Python in a non-default location and Hue is unable to locate the psycopg2 PostgreSQL database adapter.

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

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

Key: PYTHONPATH

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

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

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

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

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

None.

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

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

None.

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

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

None.

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

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

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

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

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

CDPD-43293: Unable to import Impala table using Importer

Creating Impala tables using the Hue Importer may fail.

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

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

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

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

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

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

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

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

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

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

Hue 4 web interface uses the unsafe-inline directive in its CSP header. As a result, the application server does not set the CSP header in its HTTP responses, and therefore does not benefit from the additional protection against potential cross-site scripting issues and other modern application vulnerabilities which a properly configured CSP may provide. This could lead to application vulnerability.

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

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

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

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

```
[search]
# URL of the Solr Server
solr_url=http://[***HOSTNAME***]:[***PORT***]/solr/
```

For example:

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

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

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

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

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

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

Invalid S3 URI error while accessing S3 bucket

The Hue Load Balancer merges the double slashes (//) in the S3 URI into a single slash (/) so that the URI prefix `"/filebrowser/view=S3A:/"` is changed to `"/filebrowser/view=S3A:/"`. This results in an error when you try to access the S3 buckets from the Hue File Browser through the port 8889.

The Hue web UI displays the following error: “Unknown error occurred”.

The Hue server logs record the “ValueError: Invalid S3 URI: S3A” error.

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.

You should be able to load the S3 browser from both 8888 and 8889 ports.

Alternatively, you can use the Hue server port 8888 instead of the load balancer port 8889 to resolve this issue.

Error while rerunning Oozie workflow

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

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

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

CDPD-17956, CDPD-14167: Setting idle session timeout for Hue does not work when the cluster is secured using Knox SSO

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

None

Query history displays "Result Expired" message

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

None.

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

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

Hue limitation after upgrading from CDH to CDP Private Cloud Base

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

Hue UI is blank upon login after upgrading to CDP 7.1.7 from CDH 6

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

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

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



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

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

Unsupported features

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

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

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

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

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

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

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

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

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

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



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

Known Issues in Impala

Learn about the known issues in Impala, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

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

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

The following services support DDL statements:

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

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

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

CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the `--tls_ciphersuites` startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.

Ensure that the list in the `--tls_ciphersuites` startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

Apache Jira: [IMPALA-14625](#)

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work

If you try to run `impala-shell` on RHEL 9 by setting the default python executable available in `PATH` to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, `IMPALA_PYTHON_EXECUTABLE=python3`.

CDPD-81980: Impala query may fail due to backpressure

Impala queries could fail with a dropped due to backpressure error, even when service queue and memory usage appear normal.

Increase the value of the `--mem_limit` startup option.

CDPD-42958: After upgrading the CDH 7.1.9 from CDH 6.x, under certain conditions you cannot insert data into a table

Under the following conditions, after upgrading from CDH 6.x to CDH 7.1.9 you cannot insert data into a table from Impala:

- On CDH 6.x, you created a database with Impala in a user specified HDFS location.
- Using Hive, you then created a table in the database.

Under these conditions, the database and table are stored in the user-specified HDFS directory. After upgrading, the HDFS directory of the table is read-only for Impala. Consequently, from Impala you cannot insert new data into the table because Impala does not have write permission on the HDFS directory.

Workaround: To resolve this issue, use either one of the following workarounds:

- Using the Ranger Web UI, in the policy repository `cm_hdfs`, grant the user 'impala' write permission on the directory where the table resides.
- Enter the following command to grant write permission to user 'impala' on the HDFS directory where the table resides.

```
hdfs dfs -setfacl -m default:user:impala:rwX <HDFS directory>
```

CDPD-80019: Impala fails to start due to large catalog-update topic

Impala coordinators fail to start if the catalog-update store topic exceeds 2GB, hitting the Thrift 0.16.0 message size limit (`thrift_rpc_max_message_size`). This prevents the coordinator from processing the update.

- Enable local catalog mode: `--use_local_catalog=true`
- In Cloudera Manager, go to Impala Configuration search for "Catalog Server Command Line Argument Advanced Configuration Snippet (Safety Valve)" and click 
- Add property name and value:

```
Name : --catalog_topic_mode
Value: minimal
```

- Save changes.

Apache Jira: [IMPALA-12350](#)

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 HDF
S file hdfs://nameservice1/warehouse/tablespace/managed/hive/<da
tabase>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException:
File does not exist: /warehouse/tablespace/managed/hive/<data
base>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

You must use a single Knox node to access Impala UI.

Impala api calls via 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-21828: Multiple permission assignment through grant is not working

None

Problem configuring masking on tables using Ranger

The following Knowledge Base article describes the behavior when we configure masking on tables using Ranger. This configuration works for Hive, but breaks queries in some scenarios for Impala.

For a workaround, see the following Knowledge Base article: [ERROR: "AnalysisException: No matching function with signature: mask\(FLOAT\)" when Impala jobs fail with the following error with signature: mask\(FLOAT\)](#)

IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXXXX"

IMPALA-532: Impala should tolerate bad locale settings

If the LC_* environment variables specify an unsupported locale, Impala does not start.

Add LC_ALL="C" to the environment settings for both the Impala daemon and the Statestore daemon.

Avro Scanner fails to parse some schemas

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

IMPALA-691: Process mem limit does not account for the JVM's memory usage

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the impalad daemon.

To monitor overall memory usage, use the `top` command, or add the memory figures in the Impala web UI `/memz` tab to JVM memory usage shown on the `/metrics` tab.

IMPALA-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Remove trailing semicolon from the Avro schema.

IMPALA-1652: Incorrect results with basic predicate on CHAR typed column

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the `RPAD()` function to blank-pad literals compared with `CHAR` columns to the expected length.

IMPALA-1821: Casting scenarios with invalid/inconsistent results

Using a `CAST()` function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

IMPALA-2005: A failed CTAS does not drop the table if the insert fails

If a CREATE TABLE AS SELECT operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

Drop the new table manually after a failed CREATE TABLE AS SELECT

IMPALA-2422: % escaping does not work correctly when occurs at the end in a LIKE clause

If the final character in the RHS argument of a LIKE operator is an escaped \% character, it does not match a % final character of the LHS argument.

None

IMPALA-2603: Crash: impala::Coordinator::ValidateCollectionSlots

A query could encounter a serious error if includes multiple nested levels of INNER JOIN clauses involving subqueries.

None

IMPALA-3094: Incorrect result due to constant evaluation in query with outer join

An OUTER JOIN query could omit some expected result rows due to a constant such as FALSE in another join clause. For example:

```
explain SELECT 1 FROM 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;
+---+
|
```

[illegible]

IMPALA-3509: Breakpad minidumps can be very large when the thread count is high

The size of the breakpad minidump files grows linearly with the number of threads. By default, each thread adds 8 KB to the minidump size. Minidump files could consume significant disk space when the daemons have a high number of threads.

Add `-minidump_size_limit_hint_kb=size` to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then 2 KB per thread after that.) The minidump file can still grow larger than the "hinted" size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

IMPALA-4978: Impala requires FQDN from hostname command on Kerberized clusters

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the `gethostname()` system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a Kerberized cluster.

Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname`, only returns the short name, pass the command-line flag `##hostname=FULLY_QUALIFIED_DOMAIN_NAME` in the startup options of all Impala-related daemons.

IMPALA-6671: Metadata operations block read-only operations on unrelated tables

Metadata operations that change the state of a table, like `COMPUTE STATS` or `ALTER RECOVER PARTITIONS`, may delay metadata propagation of unrelated unloaded tables triggered by statements like `DESCRIBE` or `SELECT` queries.

None

IMPALA-7072: Impala does not support Heimdal Kerberos

None

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.

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.

Known Issues in Iceberg

Learn about the known issues in Iceberg, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

There are no known issues identified in earlier releases.

Known Issues in Kafka

Learn about the known issues in Kafka, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

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

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

The following services support DDL statements:

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

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

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

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker) in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Save Changes > Restart SMM.

The offsets.topic.replication.factor property must be less than or equal to the number of live brokers

The `offsets.topic.replication.factor` broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a `GROUP_COORDINATOR_NOT_AVAILABLE` error until the cluster size meets this replication factor requirement.

None

Requests fail when sending to a nonexistent topic with `auto.create.topics.enable` set to true

The first few produce requests fail when sending to a nonexistent topic with `auto.create.topics.enable` set to true.

Increase the number of retries in the producer configuration setting `retries`.

Performance degradation when SSL Is enabled

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with `ssl.secure.random.implementation = SHA1PRNG`. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

OPSAPS-43236: Kafka garbage collection logs are written to the process directory

By default Kafka garbage collection logs are written to the agent process directory. Changing the default path for these log files is currently unsupported.

None

RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transactional method because we are in an error state
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeFailWithError(TransactionManager.java:1125)
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeAddPartition(TransactionManager.java:442)
    at org.apache.kafka.clients.producer.KafkaProducer.doSend(KafkaProducer.java:1000)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:914)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:800)
    .
    .
    .
    Caused by: org.apache.kafka.common.errors.ClusterAuthorizationException: Cluster authorization failed.
```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting `enable.idempotence` to false.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a `STRUCT` type default value.

None.

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

The Debezium Db2 Source connector does not support the evolution (updates) of schemas. In addition, schema change events are not emitted to the schema change topic if there is a change in the schema of a table that is in capture mode. For more information, see [DBZ-4990](#).

None.

CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy compression. The HDFS Stateless Sink connector is only affected if the Compression Codec or Compression Codec for Parquet properties are set to SNAPPY.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: 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

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to `required`. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to `requested` instead of `required` and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Unsupported Features

The following Kafka features are not supported in Cloudera Data Platform:

- Only Java and .Net based clients are supported. Clients developed with C, C++, Python, and other languages are currently not supported.
- The Kafka default authorizer is not supported. This includes setting ACLs and all related APIs, broker functionality, and command-line tools.
- SASL/SCRAM is only supported for delegation token based authentication. It is not supported as a standalone authentication mechanism.
- Kafka KRaft in this release of Cloudera Runtime is in technical preview and does not support the following:
 - Deployments with multiple log directories. This includes deployments that use JBOD for storage.
 - Delegation token based authentication.
 - Migrating an already running Kafka service from ZooKeeper to KRaft.
 - Atlas Integration.

Limitations

Collection of Partition Level Metrics May Cause Cloudera Manager's Performance to Degrade

If the Kafka service operates with a large number of partitions, collection of partition level metrics may cause Cloudera Manager's performance to degrade.

If you are observing performance degradation and your cluster is operating with a high number of partitions, you can choose to disable the collection of partition level metrics.



Important: If you are using SMM to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both SMM and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, SMM will not be able to display information about partitions. In addition, Cruise Control will not operate properly.

Complete the following steps to turn off the collection of partition level metrics:

1. Obtain the Kafka service name:
 - a. In Cloudera Manager, Select the Kafka service.
 - b. Select any available chart, and select Open in Chart Builder from the configuration icon drop-down.
 - c. Find \$SERVICENAME= near the top of the display.

The Kafka service name is the value of \$SERVICENAME.

2. Turn off the collection of partition level metrics:
 - a. Go to Hosts Hosts Configuration .
 - b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

Enter the following to turn off the collection of partition level metrics:

```
[KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_entity_update_enabled=false
```

Replace [KAFKA_SERVICE_NAME] with the service name of Kafka obtained in step 1. The service name should always be in lower case.

- c. Click Save Changes.

Known Issues in Knox

Learn about the known issues in Knox, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

CDPD-92941: Knox Gateway logging configuration cannot be customized through Cloudera Manager

The Knox Gateway Logging Configuration (Safety Valve) in Cloudera Manager is not applied to Knox Gateway. Individual logger configurations cannot be set or modified through the Cloudera Manager UI.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

CDPD-76435: Default HTTP client timeout for Streams Messaging Manager requests may cause connection timeouts

Some Streams Messaging Manager requests, particularly those retrieving topic data, may take up to 30 seconds to complete, even when there are no connectivity issues. Since the default gateway timeout is set to 20 seconds, this can result in connection timeouts despite the system functioning as expected.

Increase the HTTP client and socket timeouts for Streams Messaging Manager services to 60 seconds. This prevents connection timeouts when the requests take longer than the default gateway timeout.

Follow these steps to configure the required properties for the SMM API and SMM UI services:

1. In Cloudera Manager, select the Knox service.
2. Go to Configuration.
3. Find the Knox Simplified Topology Management - cdp-proxy property, and click **+** to add the following lines:

```
SMM-UI:replayBufferSize=512
SMM-UI:httpclient.connectionTimeout=1m
SMM-UI:client.socketTimeout=1m
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

4. Find the Knox Simplified Topology Management - cdp-proxy-api property, and click **+** to add the following lines:

```
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

5. Find the Knox Simplified Topology Management - cdp-proxy-token property, and click **+** to add the following lines:

```
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

6. Save your changes.
7. Refresh the cluster.

Known Issues in Kudu

Learn about the known issues in Kudu, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

CDPD-94474: Kudu Tablet Server Crash on RHEL9

Kudu tablet servers kudu-tserver crash on Red Hat Enterprise Linux 9 (RHEL9) x86_64 platforms when the libgcc package is updated to version 11.5.0-11 or newer. This crash manifests as a Segmentation Fault (SIGSEGV) or an Abort (SIGABRT) error under certain workloads.

1. Log in to Cloudera Manager as an Administrator.
2. Navigate to Cloudera Manager Clusters Kudu Configuration .
3. Navigate to the Configuration tab, and add the following value in the Tablet Server Advanced Configuration Snippet (Safety Valve) for gflagfile for the Tablet Server Default Group field:

```
--mrs_use_codegen=false
```

4. Click Save Changes.
5. Restart the Kudu Service.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

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.

Workaround: Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

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.

Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

Unsupported features

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 in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

The following known issues were identified in this release:

No new known issues in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

No known issues identified in Livy.

Known issues in Navigator Encrypt

Learn about the known issues in Navigator Encrypt, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

Failed to start navencrypt-mount.service

After installing NavEncrypt on SLES, if the command "systemctl status navencrypt-mount" fails with the error: "Failed to start navencrypt-mount.service: Unit navencrypt-mount.service failed to load: No such file or directory", the systemd files need to be reloaded.

Run the command:

```
(sudo) systemctl daemon-reload
```

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

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 Oozie

Learn about the known issues in Oozie, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

There are no known issues identified in the releases before Cloudera Runtime 7.1.9 SP1 CHF 13.

Known Issues in Ozone

Learn about the known issues in Ozone, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

CDPD-88335: Datanode decommission fails when other datanodes are offline due to invalid affinity node in Ratis replication

Datanode decommissioning fails if one or more datanodes holding Ratis replicas are offline. During decommission, the Storage Container Manager (SCM) tries to replicate Ratis container data to maintain redundancy. However, SCM's placement logic still references an affinity node that was previously part of the network topology but has since been removed due to the datanode going offline. This causes repeated placement errors and prevents the decommission process from completing successfully.

Apache Issue: [HDDS-13544](#)

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

CDPD-87429: Unexpected log lines in Ozone snapshot diff output

`ozone.client.server-defaults.validity.period.ms` configuration introduced to get server default configurations causes Replication Manager to fail when replicating data between Ozone buckets with the following error when the config value does not specify a timeunit.

```

...
java.lang.IllegalStateException: Invalid first line in diff: "25
/07/02 05:58:44 INFO Configuration.deprecation: No unit for ozon
e.client.server-defaults.validity.period.ms(3600000) assuming MI
LLISECONDS"
at com.cloudera.enterprise.distcp.OzoneCopyListingUtils.parseOz
oneDiff(OzoneCopyListingUtils.java:397)
...

```

Overwrite the `ozone.client.server-defaults.validity.period.ms` to configure the value to have time unit (ns, ms, s, m, h). For example, 3600000ms.

CDPD-71648: Snapshot creation failing with "PERMISSION_DENIED" error on a bucket created during pre-upgrade in 7.1.7 SP3

On upgrade of cluster from 7.1.7 SP3 to 7.1.9 SP1, we are not able to create snapshots on a pre-upgrade volume or bucket with "PERMISSION_DENIED" error, since in 7.1.7SP3, we were using full KRB owner name with domain but currently, we only use and search for user shortname.

Update the owner of the bucket with the user shortname and then create snapshot by running the commands:

```
ozone sh bucket update --user=[***USER_SHORT_NAME***] {volume}/
{bucket}
```



```
ozone sh snapshot create {volume}/{bucket}
{[***SNAPSHOT_NAME***]}
```

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

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

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

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

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

CDPD-64865: Intermittent timeout in TestBlockDeletion.testBlockDeletion (BlockDeletionService stuck)

If the event of SCM exiting safe mode was triggered multiple times, this might eventually cause the SCMBlockDeletingService to transition into the PAUSING state.

Check the SCM logs and ensure that the SCMBlockDeletingService is not running, and then restart the SCM.

CDPD-81569: Recon is not reflecting correct metrics for Volumes, Buckets, Keys, and Pending Key Deletions

Recon is displaying data for container and space utilisation only. The metrics for Volumes, Buckets, Keys, and Pending Key Deletions are not being updated and displayed correctly.

Apache Jira: [HDDS-12607](#)

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: 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.

The Prometheus binary is not available in CDP Private Cloud Base 7.1.9 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.

After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and if Ozone is in the Non-HA environment, an exception message is observed during the finalization of the Ozone upgrade.

During the finalization of the upgrade, a ClassNotFoundException message for org.cloudera.log4j.redactor.RedactorAppender class may be displayed. The error message can be ignored because the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

CDPD-68951: In CDP Private Cloud Base 7.1.9 CHF2 version and lower, the command `ozone sh key list <bucket_path>` displays the isFile flag in a key's metadata as false even when the key is a file. This issue is rectified in CDP Private Cloud Base 7.1.9 CHF3. However, the pre-existing (pre-upgrade) key's metadata cannot be changed.

None

CDPD-67771: When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(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.

CDPD-64398: In SCM with non-HA configuration, Secret key manager does not initialise. Hence, the startup of OM and DataNode fails because it cannot get a secret key. This key is used when security is enabled and used in Block token and container token verification while communicating between the Ozone client, OM, and Datanode.

This issue does not occur in SCM with HA configuration.

You must force exit from the safe mode for SCM. This triggers the initialization of the Secret key manager. Below are the options:

- exit safe mode manually whenever SCM is started
- disable safe mode by setting `hdds.scm.safemode.enabled=false` in safety valve for SCM configuration

Impact of disabling safe mode and the purpose of safe mode:

- Write should not fail once the cluster is out of safe mode
- Read of existing data should not fail after the cluster is out of safe mode
- Unnecessary re-replication should be avoided during cluster restart

So disabling safe mode does not have any major impact on Ozone.

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.

CDPD-52412: keyManagerImpl#listStatus exceeds the maximum RPC length.

To reduce the payload, you must increase the part size.

OPSAPS-68159: If you did not deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Error when distributing parcel to host, Ignoring non-compliant parcel manifest error is displayed after Cloudera Manager is upgraded to 7.11.3.

If you encounter the error, perform the following steps:

1. Deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.11.3.
2. Restart the cluster with a delay of 10 minutes.
3. Continue to upgrade CDH 7.1.8 to CDH 7.1.9.

OPSAPS-68159: If you did not deactivate the Ozone parcel 718.2.x on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Ozone roles during the CDH 7.1.8 upgrade to CDH 7.1.9.

If you encounter the error, perform the following steps:

1. Deactivate the Ozone parcel 718.2.x.
2. Restart the Ozone service.
3. Perform Finalize Upgrade for the Ozone service.

Step result: The Ozone roles will be displayed in green.

CDPD-60989: The packaging version for Apache Ozone points to the 1.2.0 older version. This is a version string problem and not a packaging issue. The version of the Apache Ozone binary is closest to 1.4.0.

None. This only affects the JAR versioning.

CDPD-60366: Native library loader fails when system property `native.lib.tmp.dir` is not set. It fails because the library is copied to / instead of the `cwd`.

Setting `native.lib.tmp.dir` to a certain path like `/tmp` should solve the issue. To set this system property, you must add `-Dnative.lib.tmp.dir=/tmp` to `ozone_java_opts` in Cloudera Manager configuration of the Ozone cluster.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-60598: DataNode can consume up to the JVM heap size worth of direct memory buffers under specific failure scenarios. This can lead to a DataNode crash. This is due to Netty allocating and not freeing the direct memory buffers.

Restarting DataNode will resolve this issue.

CDPD-60466: The Ozone client is missing in the Cloudera Manager's Service Monitor node. This causes Ozone Canary's health check to fail.

Ensure you install Ozone CLI on the same node as Cloudera Manager's Service Monitor.

CDPD-60012: The Ozone SCM may be stuck in safe mode after upgrade to 7.1.9. This is due to a bug in how the SCM accounts for open storage containers.

The workaround is to temporarily set the Ozone configuration `hdds.scm.safemode.threshold.pct` to a lower value like 0.90 and restart the SCM.

CDPD-59679: In some scenarios, excess UNHEALTHY replicas of EC containers may not be removed from the cluster.

None

CDPD-50116: Topology-aware reads can provide better performance by routing applications to the closest replica of a file, if available, on the same physical node or same rack. However this feature is disabled by default in CDP 7.1.9.

This feature can be enabled by setting `ozone.network.topology.aware.read` to true using Cloudera Manager safety valve.

CDPD-57165: The DataNode disk usage thread may abort due to Ratis log tailing.

Use `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory` to calculate disk usage. Configure in the `ozone-site.xml` property: `hdds.datanode.du.factory` value: `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory`

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.

CDPD-55513: Hive external table replication policy does not migrate Hive tables in HDFS storage to Ozone.

To replicate data from HDFS to Ozone, see [Migrating your data from HDFS to Ozone](#). You can use HMS Mirror to replicate the metadata.

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"
```

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_SHAREDLIB_DIR=/VAR/TMP`
3. Restart Ozone.

CDPD-49137: Ozone Manager Kerberos token expires for SCM communication and OM does not log in again.

Sometimes, OM's Kerberos token is not updated and it stops to communicate with SCM. When this occurs, writes start failing.

Restart OM or set the safety valve `hadoop.kerberos.keytab.login.autorenewal.enabled = true`

CDPD-56684: Keys get deleted when you do not have permission on volume

When a volume is deleted, it recursively deletes the buckets and keys inside it and only then deletes the volume. The volume delete ACL check is done only in the end, due to which you may end up deleting all the data inside the volume without having delete permission on the volume.



Note: There was no bucket/key permission available which allowed the user to delete them recursively.

CDPD-50610: Large file uploads are slow with OPEN and stream data approach

Hue file browser uses the append operation for large files. This API is not supported by Ozone in 7.1.9 and therefore large file uploads can be slow or timeout from the browser.

Use native Ozone client to upload large files instead of the Hue file browser.

OPSAPS-64097: Ozone service restart failed at SCM

Stopping SCM service using Cloudera Manager can sometimes timeout and need a retry. The Cloudera Manager API waits for 90 seconds which is not sufficient under certain circumstances.

Retry the shutdown using Cloudera Manager if the SCM still shows as running after a refresh of the service status.

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-filesystem-hadoop3.jar`, it is likely due to the presence of the user `hdfs` on the node prior to CDP parcel activation. As a result, the Cloudera Manager agent skips the initialization related to HDFS, leading to the non-creation of the `/var/lib/hadoop-hdfs` directory.

Create the directory `“/var/lib/hadoop-hdfs”` on all nodes followed by the deactivation and activation of the CDP parcel (deactivate and activate the Ozone parcel instead, in case Ozone parcel is used).

CDPD-50447: When SCM High Availability is enabled, each of the SCM web UIs report the host of the web UI as the leader of HA, and the other two as followers. This gives wrong information

Correct output is available by running the `ozone admin scm roles --service-id=<ID>` command.

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](#).

OPSAPS-62327: In an Ozone cluster without any gateway roles, Ozone is unable to deploy client configurations and displays the ConfigGenException error.

You must add the Ozone gateway roles to the cluster.

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-36389: The configurations `datanodes.involved.max.percentage.per.iteration` and `size.moved.max.per.iteration` are meant to limit the max number of DataNodes that is involved and maximum size that can move in an iteration. This bug causes balancer to stop an iteration when it is two DataNodes or 1 Container size (5GB) away from hitting these limits. However, these DataNodes can again be considered for balancing in the next iteration. This means the cluster is balanced after enough iterations, albeit a bit slowly. This bug is apparent in small clusters of around four DataNodes where the DataNodes can be either the source or target for a lot of moves but the iteration gets stopped when three DataNodes have been involved. It takes a higher number of iterations to eventually balance this cluster. While this is a performance issue, it does not prevent balancer from ultimately balancing the cluster.

Increase the speed for balancing by decreasing the interval between each iteration using the configuration `balancing.iteration.interval`. The value of this configuration must be greater than `hdds.datanode.du.refresh.period`. `size.moved.max.per.iteration` can be increased to allow more data to move in one iteration.

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-35141: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2> (state=08S01,code=40000) java.sql.SQLException: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2>. This issue does not occur if the source and target buckets are different in Hive queries. Currently, copying across the same bucket is only supported.

Avoid different buckets in source and target path. Ensure that `hive.exec.stagingdir` is unset or set to the same bucket as the warehouse so that staging directories are created in the same bucket as the warehouse.

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-41184: With Legacy buckets, FileSystem Optimised is not interoperating with the Ozone shell command. The cause is that the directory key entry in the DB KeyTable is stored as `"dir1/"` with trailing slash. But while performing the described operation, Ozone shell (`o3://`) is normalizing the given path and removing the trailing slash `"/"` from it. That results in `KEY_NOT_FOUND` exception.

There are three workarounds:

- Use FileSystem API to Delete the Directories rather than Shell-Command API.

- Use FSO buckets instead of Legacy Buckets. As in FSO, you can create Intermediate Directories and Delete Directories using the Ozone shell commands.
- Disable and set the configuration `ozone.om.enable.filesystem.pathsflag` to false in order to delete the directories. This is generally not a preferred workaround because the cluster must be restarted to pick up the new changes.

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-12542: Mount of Ozone filesystem with the help of FUSE fails.

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-42897: EC writes are failing with "No enough datanodes to choose" after EC replication config is set globally.

EC writes start failing when large number of pipelines are created as a result of multiple EC configs with different chunk sizes used to write keys.

If standard EC configs (that is, rs-3-2-1024k) are used to write keys, number of pipelines created per DataNode is limited to five and this issue is not seen with standard EC configs.

The recommendation is not to create too many random chunk sizes. It is configurable because, users can decide based on their workload. But do not to have separate chunk sizes for each file.

CDPD-41539: The "No such file or directory" imessage is displayed when EC file is read using older ofs client.

You must upgrade the client before trying to read the key: vol1/ecbuck1/1GB_ec".

CDPD-40560: Filesystem Operations using Hadoop S3a connector on a FILE_SYSTEM_OPTIMIZED bucket fails. org.apache.hadoop.ozone.om.exceptions.OMException: Unable to get file status: volume: s3v bucket: fso key: test/

Do not run Hadoop S3a commands on a FILE_SYSTEM_OPTIMIZED bucket. Use OBJECT_STORE bucket layouts.

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.

CDPD-43432: Ozone Service in fault state in DataNode - Long Running setup.

Upgrade RocksDB to the latest version.

OPSAPS-63999: In the newly installed cluster, the Finish upgrade option is clickable.

None

OPSAPS-64648: Failed to start Ozone node using Cloudera Manager if default log path /var/log/hadoop-ozone does not exist. If this path does not exist, any Ozone nodes(for example SCM or DataNode) restart will fail.

Run the following command `sudo -u hdfs mkdir -p /var/log/hadoop-ozone` or replace hdfs with the user Ozone roles that are running.

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.

Known Issues in Ranger

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

Access issues with users having similar name as admin or keyadmin users

Users with USER role that have a similar name as admin or keyadmin users are able to query those admin and keyadmin users.

None.

UnsupportedClassVersionError

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.

CDPD-68806: The Revoke operation for users belonging to a group or role permission does not function as expected.

List command is listing all the tables even when the user permission is revoked. And also the command does not add any deny policy to Ranger for that specific user.

For example:

user_1 is the part of group_1

user_2 is the part of role_1

Create a Ranger policy for group_1 and role_1 to access all the resources in HBase.

Try to revoke the permission for user_1 and user_2 using HBase shell global revoke command:

```
revoke 'user_1'
```

```
revoke 'user_2'
```

But there is no deny policy added for that user. And access is still granted to user_1 and user_2.

This behavior is currently not supported in HBase shell. Must be handled manually using the Ranger policy change.

CDPD-68739: The revoke command does not work when using the HBase shell

While using the HBase shell, running the revoke command does not cancel the user permission. Users are able to perform actions even after running the revoke command.

For example:

When a user is granted permission to a particular table using the HBase grant command.

```
hbase:003:0> grant 'user_1', 'RWCA', 'some_abc_table'
```

And then execute the global revoke command for that user:

```
hbase:005:0> revoke 'user_1'
```

Permission for that user 'user_1' for the table 'some_abc_table' not got revoked.

None.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-58860: After upgrading CDP-7.1.8 to CDP-7.1.9, cdp-proxy token missing from Knox-Ranger policy.

As part of [OPSAPS-67480](#) in 719 default ranger policy is added from cdp-proxy-token topology, so that after a new installation of CDP-7.1.9, the Knox-Ranger policy includes cdp-proxy-token. However, upgrades do not add cdp-proxy-token to cm_knox policies automatically.

Manually add cdp-proxy-token to the Knox policy, using Ranger Admin Web UI.

1. Log in to Cloudera ManagerRangerRanger Admin Web UI, as a Ranger administrator.
2. On Ranger Admin Web UIService ManagerResourceKnox, click cm_knox.
3. In Knox Policies, open the CDP Proxy UI, API and Token policy.
4. In Knox Topology*, add cdp-proxy-token.
5. Click Save.
6. Restart Ranger.

CDPD-58704: hadoop roll key / key delete command shows operation failed error when one KMS host is down, even when operation succeeds

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered kms instances. if one KMS instance is stopped (not deleted), client gets 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.

```
hbase:001:0> revoke 'hrt_11'
ERROR: org.apache.hadoop.hbase.coprocessor.CoprocessorException: HTTP 400 Error: processSecureRevokeRequest processing failed
    at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor.preRevoke(RangerAuthorizationCoprocesor.java:1309)
    at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor.preRevoke(RangerAuthorizationCoprocesor.java:1128)
    at org.apache.hadoop.hbase.master.MasterCoprocesorHost$162.call(MasterCoprocesorHost.java:1857)
    at org.apache.hadoop.hbase.master.MasterCoprocesorHost$162.call(MasterCoprocesorHost.java:1854)
    at org.apache.hadoop.hbase.coprocessor.CoprocessorHost$0bserverOperationWithoutResult.callObserver(CoprocessorHost.java:558)
    at org.apache.hadoop.hbase.coprocessor.CoprocessorHost.executeOperation(CoprocessorHost.java:631)
    at org.apache.hadoop.hbase.master.MasterCoprocesorHost.preRevoke(MasterCoprocesorHost.java:1854)
    at org.apache.hadoop.hbase.master.MasterRpcServices.revoke(MasterRpcServices.java:2740)
    at org.apache.hadoop.hbase.shaded.protobuf.generated.MasterProtos$MasterService$2.callBlockingMethod(MasterProtos.java)
    at org.apache.hadoop.hbase.ipc.RpcServer.call(RpcServer.java:387)
    at org.apache.hadoop.hbase.ipc.CallRunner.run(CallRunner.java:139)
    at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:369)
    at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:349)
```

None

CDPD-56741: Improvement in log message when jwtauth not used

The following exception is printed at startup only and it is not cluttering logs:

```
2023-05-30 06:18:40,127 ERROR org.apache.ranger.rms.security.RMSJwtAuthFilter:
quasar-pibgzl-1.quasar-pibgzl.root.hwx.site-startStop-1]:
Failed to initialize Ranger RMS JWT Auth Filter.
java.lang.Exception: RangerJwtAuthHandler:
Mandatory configs ('jwks.provider-url' & 'jwt.public-key')
are missing, must provide atleast one.
    at org.apache.ranger.authz.handler.jwt.RangerJwtAuthHandler.initialize(RangerJwtAuthHandler.java:84)
    ~[ranger-authn-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
```

```

        at org.apache.ranger.rms.security.RMSJwtAuthFilter.initialize(RMSJwtAuthFilter.java:77)
        ~[ranger-rms-common-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
        at jdk.internal.reflect.NativeMethodAccessorImpl.invoke0(Native Method) ~[?:?]
        at jdk.internal.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl

```

This exception is logged, since mandatory JWT auth filter configs are not provided (expected for some environments like y-cloud) Even though JWT auth filter is failed to initialised, it falls back to kerberos auth, hence no impact from auth perspective

NA

CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup

This is a warning log printed in catalina.out file when Ranger RMS server is initialised. Following exception is observed only in Oracle 19 setup. FileNotFoundException: /usr/share/java/oraclepki.jar

NA

CDPD-48975: Ranger KMS KTS to KMS DB migration : keys with the same name but different case are not migrated

KMS keys are not case sensitive

No workarounds. Such keys combination are very rare and migration doc was updated to check such keys before starting the migration.

CDPD-42598: Kafka policy creation allowed with incorrect permissions.

When creating a Kafka policy from the UI, the permissions "Idempotent write" and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource

None.

CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload:

```
{"resourceName": "classification", "userInput": "", "resources": {"classification": []}}
```

expectation is to return all the classifications . But the response has only "classification" Happens similarly for entity-label , entity-business-metadata.

None.

CDPD-40734: User allowed to insert data into a hive table when there is a deny policy on a table column.

A user is allowed to enter data into a table even if there is a deny policy present on one of the table columns.

Test scenario details:

```

Policy setup :-
policy 1 :- all access policy for hrt_qa, hive and impala
users
resources - database - * , table - *, column - *
users : hrt_qa, hive, impala
access - all access allowed
policy 2 :- policy on test_1.table_1 for hrt_5
users : hrt_5
resources : database - test_1, table - table_1, column - *
access :- all access allowed

```

```

policy 3 :- deny policy on test_1.table_1.c0 for hrt_5
users : hrt_5
resources : database - test_1, table - table_1, column -
c0
access - all access denied
data setup :-
database - test_1
table - table_1(c0 int, c1 int)

```

The user is able to insert data into the table.

None.

Known Issues in Schema Registry

Learn about the known issues in Schema Registry, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

CDPD-40380: Authorization checking issue when Kerberos is disabled

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the `ranger.plugin.schema-registry.service.name` field.
3. Replace `GENERATED_RANGER_SERVICE_NAME` with the actual name of the service.
4. Restart the Schema Registry service.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a `STRUCT` type default value.

None.

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

Schema Registry uses a schema fingerprinting mechanism to differentiate between schemas. However, fingerprinting does not take into consideration the schema attributes of the field type. As a result, if you have two schemas where the only difference is that one has type attributes defined and the other does not, they will be considered identical by Schema Registry. For example, the following schemas are considered identical:

```

#Schema V1
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"local_timestamp_micros_long","type":"long"}]}

#Schema V2
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"local_timestamp_micros_long","type":{"type":"long","logicalType":"local-timestamp-micros"}]}

```

Notice that the only difference is that in the second schema, the `local_timestamp_micros_long` field has a logical type specified. In cases like this, the new version of the schema is not created, the initial version is used. This is true even if the data that is being produced has a new schema version. The ID of the first schema version is used and is put in the serialized record. The new schema version is not created.

This issue is common when using change data capture (CDC) connectors like the Debezium connectors. This is because CDC connectors create schemas with the logical type decimal based on the column type in the database schema. For example:

```
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [{"name": "database_column", "type": {"type": "bytes", "logicalType": "decimal", "precision": 64, "scale": 0}}]}
```

If the database schema changes (for example, the column type), it is possible that only scale changes, which is a schema attribute.

```
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [{"name": "database_column", "type": {"type": "bytes", "logicalType": "decimal", "precision": 64, "scale": 1}}]}
```

In this case, even though scale changed to 1, the first version of the schema is used where scale is 0. As a result, the data is consumed with the wrong scale.

Avoid using logical types or other attributes. Alternatively, ensure that there are no changes in the logical types or other attributes between schema versions.

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

Schema Registry does not start if the address specified in the Load Balancer Address Ranger property does not end with a trailing slash (/).

Set the value of the RANGER_REST_URL Schema Registry environment variable to an address that includes a trailing slash.

1. In Cloudera Manager, select the Schema Registry service.
2. Go to Configuration.
3. Find the Schema Registry Server Environment Advanced Configuration Snippet (Safety Valve) property and add the following:

```
Key: RANGER_REST_URL
Value: [***RANGER REST API URL***]
```

Replace [***RANGER REST API URL***] with an address that can be used by Schema Registry to access Ranger. Ensure that the address ends with a trailing slash. For example: `http s://ranger-1.cloudera.com:6182/`

4. Restart the Schema Registry service.

OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

1. Access the Ranger Console (Ranger Admin web UI).
2. Click the `cm_atlas` resource-based service.
3. Add the `schemaregistry` user to the `all - *` policies.
4. Click `Manage Service Edit Service`.
5. Add the `schemaregistry` user to the `default.policy.users` property.

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

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect

instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Known Issues in Spark

Learn about the known issues in Spark the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

The following known issues were identified in this release:

No new known issues in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the RuntimeWarning: Failed to add file message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

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

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

The following services support DDL statements:

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

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

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

CDPD-67517: Spark3 tests are failing if /tmp is mounted as NOEXEC.

Map the tmpdir to a writable path in spark3-conf/spark-defaults.conf using the following steps:

1. Go to the Cloudera Manager.
2. Go to the Spark service.
3. Click the Configuration tab.
4. Select Scope Gateway .
5. Select Category Advanced .
6. Locate the Spark Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-defaults.conf_client_config_safety_valve property.
7. Map the tmpdir to a writable path:

```
spark.driver.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
spark.executor.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
```

8. Enter a Reason for change, and then click Save Changes to commit the changes.
9. Deploy the client configuration.

CDPD-23817: In the upgraded Cluster, the permission of /tmp/spark is restricted due to the HDP configuration `hive.exec.scratchdir=/tmp/spark`.

If you are using the /tmp/spark directory in the CDP cluster, you must provide the required additional Policies/ACL permissions.

CDPD-22670 and CDPD-23103: There are two configurations in Spark, "Atlas dependency" and "spark_lineage_enabled", which are conflicted. The issue is when Atlas dependency is turned off but spark_lineage_enabled is turned on.

Run Spark application, Spark will log some error message and cannot continue. That can be restored by correcting the configurations and restarting Spark component with distributing client configurations.

CDPD-23007: Mismatch in the Spark Default DB Location. In HDP 3.1.5, hive_db entities have one attribute - 'location' which is configured to the '/managed' path. In fresh install of CDP 7.1.7, hive_db entities now have 2 attributes 'location' configured to '/external' path and 'managedLocation' configured to '/managed' path. In the AM2CM migration (HDP 3.1.5 -> CDP 7.1.7), the 'location' attribute from hive_db entities in HDP 3.1.5 comes unaltered to CDP 7.1.7 and hence maps to '/managed' path.

This issue arises only if you are upgrading from HDP 3.1.5 to CDP 7.1.7. If you are performing a fresh install of CDP 7.1.7, you can ignore this issue.

None

CDPD-217: The Apache Spark connector is not supported

The old *Apache Spark - Apache HBase Connector* (shc) is not supported in CDP releases.

Use the new HBase-Spark connector shipped in CDP release.

CDPD-3038: Launching pyspark displays several HiveConf warning messages

When pyspark starts, several Hive configuration warning messages are displayed, similar to the following:

```
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of name hive.meta
store.runworker.in does not exist
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of
name hive.masking.algo does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of
name hive.metastore.runworker.in does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of na
me hive.masking.algo does not exist
```

These errors can be safely ignored.

Known Issues in Sqoop

Learn about the known issues in Sqoop, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version

If you have upgraded to CDP PvC Base 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to CDP PvC Base 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see [Troubleshooting Apache Sqoop issues](#).

CDPD-44431: Using direct mode causes problems

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the `--direct` option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the `sqoop.enable.deprecated.direct` property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through `-Dsqoop.enable.deprecated.direct=true`.

CDPD-3089: Avro, S3, and HCat do not work together properly

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

Parquet columns inadvertently renamed

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.

PARQUET-99: Importing Parquet files might cause out-of-memory (OOM) errors

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.

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 in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13**OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager**

Cloudera Manager does not display a Log Files menu for SMM UI role (and SMM UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by SMM UI is not supported by Cloudera Manager.

View the SMM UI logs on the host.

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in SMM safety valve in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server
Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Add the
following value for bootstrap servers>Save Changes > Restart SMM.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

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

KNOX does not send neither the Authentication header nor the hadoop.auth cookie. Because of this, the SMM UI returns an HTTP 401 response and sets the www-authenticate: Negotiate header. After this, KNOX still does not send the Authentication header. This results in the SMM UI being inaccessible through Knox.

You can access the SMM UI directly in Cloudera Manager at Clusters SMM Streams Messaging Manager Web UI and log in using the Kerberos username and password.

CDPD-82560: Streams Messaging Manager known unprotected endpoint

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

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

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 in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

CDPD-11079: Blacklisted topics appear in the list of replicated topics

If a topic was originally replicated but was later excluded for replication, 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, this topic will be included in the response. Additionally, the excluded topic will also be visible in the SMM UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

CDPD-30275: SRM may automatically re-create deleted topics on target clusters

If auto.create.topics.enable is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with srm-control. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1]
```

2. Wait until SRM is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

Limitations

SRM cannot replicate Ranger authorization policies to or from Kafka clusters

Due to a limitation in the Kafka-Ranger plugin, SRM cannot replicate Ranger policies to or from clusters that are configured to use Ranger for authorization. If you are using SRM to replicate data to or from a cluster that uses Ranger, disable authorization policy synchronization in SRM. This can be achieved by clearing the Sync Topic Acls Enabled (sync.topic.acls.enabled) checkbox.

SRM cannot ensure the exactly-once semantics of transactional source topics

SRM data replication uses at-least-once guarantees, and as a result cannot ensure the exactly-once semantics (EOS) of transactional topics in the backup/target cluster.



Note: Even though EOS is not guaranteed, you can still replicate the data of a transactional source, but you must set isolation.level to read_committed for SRM's internal consumers. This can be done by adding [***SOURCE CLUSTER ALIAS***]->[***TARGET CLUSTER ALIAS***].consumer.isolation.level=read_committed to the Streams Replication Manager's Replication Configs SRM service property in Cloudera Manager. The isolation.level property can be set on a global connector or replication level. For example:

```
#Global connector level
connectors.consumer.isolation.level
=read_committed
#Replication level
uswest->useast.consumer.isolation.level=read_committed
```

SRM checkpointing is not supported for transactional source topics

SRM does not correctly translate checkpoints (committed consumer group offsets) for transactional topics. Checkpointing assumes that the offset mapping function is always increasing, but with transactional source topics this is violated. Transactional topics have control messages in them, which take up an offset in the log, but they are never returned on the consumer API. This causes the mappings to decrease, causing issues in the checkpointing feature. As a result of this limitation, failover operations for transactional topics is not possible.

Known Issues in YARN 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 in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

None.

COMPX-4644: Queue capacity rounding problem when configuration is initially set via YARN

When setting the capacity scheduler configuration through the YARN/Cloudera Manager configuration, there may be capacity values that use multiple decimal places. This results in rounding/floating point precision discrepancies in the UI when trying to validate that all sibling capacities equal 100%. The UI looks like all the numbers add up to 100, but the validation still displays an error and does not allow to save the capacities. It is also observed that the capacity is being calculated as, for example, 99.9999999991 in the backend.

- Create queues within the UI, or
- Ensure that capacities configured through the Capacity Scheduler safety valve do not have more than one decimal place.

20202 Database migration after enabling opt-in migration

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to CDP 7.1.9, forcing migration from H2 to PostgreSQL database.
- Upgrade to CDP 7.1.9 CHF2, moving back to H2 database.
- Upgrade to CDP 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version CDP 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

CDPD-56559: MapReduce jobs can intermittently fail during a rolling upgrade.

During a rolling upgrade between CDP versions 7.1.8 and 7.1.9, MapReduce jobs may fail with message, `RuntimeException: native snappy library not available`. Although the native Snappy compression library is not loaded, a checkmark displays to indicate that the Snappy compression library is loading for NodeManagers that are pending upgrades. This causes the MapReduce jobs that are associated with the NodeManagers to fail. After the upgrade, the jobs work as expected. This issue only impacts rolling upgrades from before CDP 7.1.9 to a higher version.

None.

COMPX-12021 Queue Manager configurations on Scheduler Configuration page are not working

When setting the following properties on the YARN Queue Manager UI, the properties are set in the `capacity-scheduler.xml` file which does not have any effect on YARN. The properties need to be set in the `yarn-site.xml` file, which does not happen when you set them through YARN Queue Manager.

- "Maximum Application Priority" – `yarn.cluster.max-application-priority`

- "Enable Monitoring Policies" – "yarn.resourcemanager.scheduler.monitor.enable"
- "Monitoring Policies" – "yarn.resourcemanager.scheduler.monitor.policies"



Note: You can also set this property on the YARN configuration page in Cloudera Manager as "Capacity Scheduler Preemption".

- "Preemption: Observe Only" – "yarn.resourcemanager.monitor.capacity.preemption.observe_only"
- "Preemption: Monitoring Interval (ms)" – "yarn.resourcemanager.monitor.capacity.preemption.monitoring_interval"
- "Preemption: Maximum Wait Before Kill (ms)" – "yarn.resourcemanager.monitor.capacity.preemption.max_wait_before_kill"
- "Preemption: Total Resources Per Round" – "yarn.resourcemanager.monitor.capacity.preemption.total_preemption_per_round"
- "Preemption: Over Capacity Tolerance" – "yarn.resourcemanager.monitor.capacity.preemption.max_ignored_over_capacity"
- "Preemption: Maximum Termination Factor" – "yarn.resourcemanager.monitor.capacity.preemption.natural_termination_factor"
- "Enable Intra Queue Preemption" – "yarn.resourcemanager.monitor.capacity.preemption.intra-queue-preemption.enabled"

1. In Cloudera Manager, select the YARN service.
2. Click the Configuration tab.
3. Search for yarn-site.xml.
4. Under YARN Service Advanced Configuration Snippet (Safety Valve) for yarn-site.xml, add the corresponding parameter and value you need.
5. Click Save Changes.
6. Restart the YARN services.

COMPX-13177: QueueManager webapp requests fail with 'HTTP ERROR 400

java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA'

Products:

- Cloudera Manager for CDP Private
- Cloud Base Cloudera Manager for CDP Public Cloud

Context:

- Centos 7.8 and Redhat 7.8 operating systems, when FIPS support is enabled.

Problem:

- When attempting to display the Yarn Queue Manager interface, Cloudera Manager displays an error: "HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA".

1. Edit the file [/etc/default/cloudera-scm-server]
2. Around line 28, modify the line that starts with

```
#export CMF_OVERRIDE_TLS_CIPHERS=...
```

3. Remove the comment mark #.
4. Remove all ciphers with "3DES" in the name.
5. Save the file.
6. Restart the Cloudera Manager Server service.

Third-party applications do not launch if MapReduce framework path is not included in the client configuration

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the `mapreduce.application.framework.path` property is set to the appropriate value, but third-party applications with their own configurations does not launch.

Set the `mapreduce.application.framework.path` property to the appropriate configuration for third-party applications.

JobHistory URL mismatch after server relocation

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated the `mapred-site.xml` file that references the correct JobHistory Server.

YARN cannot start if Kerberos principal name is changed

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN does not start. In such cases, the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the `znode` and restart the YARN service.
- Use the `reset ZK ACLs` command. This also sets the `znodes` below `/rmstore/ZKRMStateRoot` to `world:anyone:cdw` which is less secure.

Queue Manager does not open on using a custom user with a default Kerberos principal

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager `process_username` property matches the YARN `process_username` property.

Known Issues in Zeppelin

Learn about the known issues in Zeppelin the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

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.

BUG-125263: Zeppelin service move fails on clusters upgraded from HDP3.1.5

Resolve the circular symlink issue on the Zeppelin node by linking the `conf` directory to a new directory under `/etc/zeppelin`:

- `# mkdir -p /etc/zeppelin/<version>/0`
- `# rm /usr/hdp/<version>/zeppelin/conf`
- `# ln -s /etc/zeppelin/<version>/0 /usr/hdp/<version>/zeppelin/conf`

Where `version` is the HDP version. For example, `7.1.x-yyy`. Restart the Zeppelin server in Ambari.

CDPD-3090: Due to a configuration typo, functionality involving notebook repositories does not work

Due to a missing closing brace, access to the notebook repositories API is blocked by default.

From the CDP Management Console, go to Cloudera Manager for the cluster running Zeppelin. On the Zeppelin configuration page (Zeppelin serviceConfiguration), enter shiro urls in the Search field, and then add the missing closing brace to the notebook-repositories URL, as follows:

```
/api/notebook-repositories/** = authc, roles[{{zeppelin_admin_group}}]
```

Click Save Changes, and restart the Zeppelin service.

CDPD-2406: Logout button does not work

Clicking the Logout button in the Zeppelin UI logs you out, but then immediately logs you back in using SSO.

Close the browser.

OPSAPS-59802: Zeppelin and Livy roles should be co-located on the same host.

When installing or upgrading to CDP Private Cloud Base, you must co-locate all Zeppelin and Livy roles on the same cluster host due to an issue with certificate generation.

Known Issues in ZooKeeper

Learn about the known issues in ZooKeeper the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 13.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 13

There are no new known issues identified in the previous releases.

Cumulative hotfix 12

Know more about the cumulative hotfix 12 for CDP 7.1.9 SP1. This cumulative hotfix was released on October 21, 2025.



Important:

The minimum required Cloudera Manager version for Cloudera Runtime 7.1.9 SP1 CHF 12 is Cloudera Manager 7.13.1 CHF4 (7.13.1.400).

Cloudera Runtime 7.1.9.1064-1 (Cumulative Hotfix 12) download URL:

Parcel Repository Location

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1064/parcels/
```

What's new in Cloudera Runtime 7.1.9 SP1 CHF 12

Understand the functionalities and improvements to features of components in Cloudera Runtime 7.1.9 SP1 CHF 12.

What's new in Ozone

Increased default value of ozone.om.ratis.segment.size and ozone.scm.ha.ratis.segment.size

Previously, the default value of the ozone.om.ratis.segment.size and ozone.scm.ha.ratis.segment.size configuration parameters was 4 MB. Now, the default value is increased to 64 MB.

- The ozone.om.ratis.segment.size configuration parameter sets the size of the raft segment used by Apache Ratis on Ozone Manager.

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

Increased default value of `hdds.datanode.volume.min.free.space`

Previously, the default value of the `hdds.datanode.volume.min.free.space` minimum free space configuration was 5 GB. Now, the default value is increased to 20 GB.

Fixed issues in 7.1.9 SP1 CHF 12

Know more about the list of fixes that are shipped for CDP Private Cloud Base version 7.1.9 SP1 CHF 12.



Note: Contact Cloudera Support for questions related to any specific hotfixes. The version and build identifier for this CDH Parcel is 7.1.9-1.cdh7.1.9.p1064.71379092.

CDPD-82979: Added Cloudera-specific configuration for MapPartitionsRDD default determinacy

To avoid failures during error handling, a Cloudera-specific configuration was added to change the default determinacy level for any MapPartitionsRDD.

CDPD-91213: Non-deterministic SQL expressions failed to set indeterminate map stage output level

Previously, a bug in Spark related to handling non-deterministic data types could cause data corruption or loss during task failures and retries. This issue is now fixed.

Apache Jira: [SPARK-51272](#)

CDPD-89515: Fixed assertion error in Spark submit application

A missing commons-configuration dependency in the spark-atlas-connector assembly jar failed caused an assertion error, as the spark-submit application did not return an application ID. This issue is now fixed.

CDPD-87670: [Atlas UI] Apache Atlas Glossary becomes unresponsive when the page size is set to 50

Previously, when the page limit was set to 50 in **Glossary**, the user interface was unresponsive due to an incorrect API call with an error message, when trying to show more than 25 associated entities for a glossary item. This issue is now fixed by updating the API function logic, and now more than 50 associated glossary items can be shown.

Apache Jira: [ATLAS-5067](#)

CDPD-88447: Prevent unchecked scheduling of flush operations under memory pressure

Previously, in some clusters, tablet servers (tservers) maintained high memory usage up to 60%–80%, causing the maintenance manager to limit operations to WAL garbage collection and MRS/DRS flushes. This led to degraded performance and potential OOM failures.

The issue is now resolved with a configurable parameter that allows the maintenance manager to continue additional operations under moderate memory pressure. It dynamically adjusts MRS/DRS flush probability based on memory usage thresholds, improving resource management and preventing performance issues.

Apache Jira: [KUDU-3407](#)

CDPD-75829: Memory leak when using Ozone FS filesystem through Hadoop FileContext API

Previously, memory used to leak when the Hadoop `FileContext` API was used. The Hadoop `FileContext` API internally uses the `AbstractFileSystem` API that is implemented by the `org.apache.hadoop.fs.ozone.RootedOzFs` class when using the Ozone FS filesystem. The `org.apache.hadoop.fs.ozone.RootedOzFs` class internally creates a new `RootedOzoneFileSystem` instance that never closes and might cause a memory leak. This issue is now fixed.

Apache Jira: [HDDS-11635](#)

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

Previously, the Ozone Manager crashed when the follower was lagging and executed the `purgeKeys` or `snapshotMoveDeletedKeys` task for the Snapshot in a single transaction. This issue is now fixed.

Apache Jira: [HDDS-11152](#)

CDPD-73318: Datanode cannot report for a long time

Previously, some datanodes could not report for a long time and their status was DEAD or STALE. This issue is now fixed.

Apache Jira: [HDDS-11331](#)

CDPD-89911: High hash collision could slow down table renames

Previously, performance was low when renaming tables with a large number of columns. This issue is addressed by optimizing the hashing function for PartColNameInfo component.

Apache Jira: [HIVE-29165](#)

CDPD-90263: Performance degradation with LIMIT clause in vectorized queries

Some vectorized queries that included a LIMIT clause would process all rows instead of stopping after the limit was reached. This issue is now resolved by modifying the vectorized code to correctly check if child operators are done, allowing the query to terminate early and significantly improving performance.

Apache Jira: [HIVE-24930](#)

CDPD-90821: Bucket map join query hangs due to auto-parallelism changes

Queries using a bucket map join could hang when auto-reducer parallelism changed the number of tasks for a source vertex. This issue is now resolved by temporarily disabling when auto-reducer parallelism for any vertices that are the source of a bucket map join.

Apache Jira: [HIVE-27078](#)

CDPD-91157: Incorrect results from bucket map join on decimal bucketed columns

Queries using a bucket map join on tables with decimal bucketed columns could return incorrect results.

The issue is now resolved by ensuring that the bucket map join correctly gathers bucket expression indices, allowing for a complete hash table and accurate results.

Apache Jira: [HIVE-27267](#)

CDPD-89389: Major compaction could lead to data loss

A major compaction could cause data loss when a table was overwritten and then new data was inserted.

This issue is now resolved by ensuring that the compactor correctly handles these bucket mismatches, so all base files are included in the major compaction, preventing data loss.

Apache Jira: [HIVE-28700](#)

CDPD-89791: Incorrect COUNT(DISTINCT) results in subqueries

Queries using the COUNT(DISTINCT) SQL function within a subquery were returning incorrect results.

This issue is now resolved by ensuring that the optimizer uses the most current plan, fixing the incorrect count.

Apache Jira: [HIVE-27084](#)

CDPD-89852/ CDPD-89853: Crash when casting timestamp strings with timezone offsets to DATE

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

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

Apache Jira: [IMPALA-14383](#)

CDPD-89964/CDPD-88861: SHA2 function fails in FIPS mode

When using the SHA2 function with a bit length of 384 in a FIPS-compliant environment, Impala crashed.

This issue is now resolved by updating the function to ensure compatibility with FIPS mode, preventing the crash.

Apache Jira: [IMPALA-14385](#)

CDPD-90828: Revert HDDS-8923 since it is using a lot of memory

HDDS-8923 added CacheMetrics to XceiverClientManager and created a unique name per instance. This process could use a lot of memory. This issue is now fixed, as HDDS-8923 is now reverted.

COMPX-22905: Null Pointer Exception in FifoCandidatesSelector during preemption check in Scheduling Monitor

A Null Pointer Exception occurred during preemption because childless parent queues were being misidentified as leaf queues. This issue is now resolved by adding an explicit check to ensure that only a true instance of the Abstract Leaf Queues are processed.

CDPD-90931: Emit offset syncs more often than offset.lag.max for low-throughput or finite partitions

Previously, Streams Replication Manager (SRM) only emitted offset syncs based on message count. This could cause steady offset translation lag (up to the value of offset.lag.max) for finite-length or low-throughput topics.

This issue is now resolved by adding time-based offset sync emission. This ensures offset syncs are emitted regularly, even for slow or stopped topics

Apache Jira: [KAFKA-15906](#)

COMPX-22434: Queue Manager failed to recreate DQS rule after exiting mixed mode

Previously, Queue Manager failed to recreate the Dynamic Queue Schedules (DQS) rule after exiting YARN mixed mode. This issue is now resolved.

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

CVE-2025-58754	Axios
CVE-2025-48989	Apache Tomcat

Behavioral changes in Cloudera Runtime 7.1.9 SP1 CHF 12

Know more about the behavioral changes in cumulative hotfix 12 for 7.1.9 SP1.

CDPD-88944: New configuration available to disable the Partition Management Task

Previous behavior:

Disable the PartitionManagementTask, which put a heavy load on the Hive Metastore when managing a large number of tables and partitions. Users had a tedious workaround by providing a pattern through metastore.partition.management.database.pattern or metastore.partition.management.table.pattern.

New behavior:

You can now set the metastore.partition.management.task.frequency configuration to 0 to disable the partition management task cluster-wide.

Apache Jira: [HIVE-25324](#)

Known issues in 7.1.9 SP1 CHF 12

The following section lists the known issues and limitations, the areas of impact, and workaround for each Runtime component.

Known Issues in Atlas

Learn about the known issues in Atlas the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

CDPD-93307: Hive import logs are truncated during bulk import

When running the import-hive.sh script to perform bulk Hive metadata imports, the logging process fails to retain the expected history. The import-hive.log file rotates prematurely when it reaches 10 MB and retains only one backup file (import-hive.log.1), instead of the expected 100 MB file size and 20 backup files.

Supply a custom log4j configuration file with the correct structure by using the ATLAS_OPTS environment variable before running the import script:

```
<?xml version="1.0" encoding="UTF-8" ?>
<!DOCTYPE log4j:configuration SYSTEM "log4j.dtd">
<log4j:configuration xmlns:log4j="http://jakarta.apache.org/log4j/">
  <appender name="console" class="org.apache.log4j.ConsoleAppender">
    <param name="Target" value="System.out"/>
    <layout class="org.apache.log4j.PatternLayout">
      <param name="ConversionPattern" value="%d %-5p - [%t:
%x] ~ %m (%C{1}:%L)%n"/>
    </layout>
  </appender>
  <appender name="FILE" class="org.apache.log4j.RollingFileAppender">
    <param name="File" value="${atlas.log.dir}/${atlas.log.file}"/>
    <param name="Append" value="true"/>
    <param name="maxFileSize" value="100MB"/>
    <param name="maxBackupIndex" value="20"/>
    <layout class="org.apache.log4j.PatternLayout">
      <param name="ConversionPattern" value="%d %-5p - [%t:
%x] ~ %m (%C{1}:%L)%n"/>
    </layout>
  </appender>
  <logger name="org.apache.atlas" additivity="false">
    <level value="info"/>
    <appender-ref ref="FILE"/>
  </logger>
  <logger name="org.apache.kafka.common.config.AbstractConfig" additivity="false">
    <level value="error"/>
    <appender-ref ref="FILE"/>
  </logger>
  <root>
    <priority value="warn"/>
    <appender-ref ref="FILE"/>
  </root>
</log4j:configuration>
```

Export the ATLAS_OPTS variable to point to your custom file before running the import script:

```
export ATLAS_OPTS="-Dlog4j.configuration=file:/path/to/custom-hive-import-log4j.xml"
```

```
./import-hive.sh
```

CDPD-87670: [Atlas UI] Apache Atlas Glossary becomes unresponsive when the page size is set to 50

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

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-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-6565: Issue with ddlQueries in Atlas for Hive/Impala tables created in default/non-default db using DAS/HUE respectively

ddlQueries are not created for Impala origin table. ddlQueries created for Impala/Hive CTAS tables have difference in names as compared to same operation for Hive table. The Impala table has the query in the ddlQueries entity name.

None

CDPD-55301: ddlQueries and ALTERNATIVE_* lineage missing for Spark tables created through spark3-shell

The ddlQueries and ALTERNATIVE_* lineage missing for Spark tables created using spark3-shell.

None

CDPD-56085: 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

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*.

None

CDPD-58554: Discard audits of specific classification, label, and business metadata

Support to control audits for specific classification, label, and business metadata is not present in the Custom Audit Filters feature.

None

CDPD-40346: The ddlQueries and ALTERNATIVE_ADDCOLS lineage missing for Impala tables.

The ALTERNATIVE_ADDCOLS lineage has some issue when an Impala table is altered and the corresponding lineage is not created.

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

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

None

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

Adding Japanese labels or user defined properties results in the error message: "Invalid label: ###, label should contain alphanumeric characters, _ or -"

None

CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties

Entities, such as `hdfs_path`, are not returned by the search when their indexable properties are searched using partial search terms made of Japanese or Chinese characters. Only exact matches return results when searching the indexed properties made of Japanese or Chinese characters.

None

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

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

None.

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

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

None.

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

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

None.

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

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

None.

CDPD-43772: Performance issues with Atlas service

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

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

OPSAPS-67783: During rolling upgrade one among two Atlas server failed to start but Cloudera Manager considered as success

Cloudera Manager marks the Execute command `Start on service Atlas-1` as a success even when Atlas service had failed to start successfully. In such cases, Atlas logs give the exact reason for Atlas start-up failure.

None

CDPD-19358: "IsIndexable" and "isOptional" value of a typedef's attribute is modified post migration.

In HDP-265, `falcon_feed_creation` has "stored-in" attribute which has `IsIndexable` value set to `True`.

Post migration, "stored-in" is moved to `relationshipAttributeDefs` and has `IsIndexable` value set to `False`.

Similarly, in HDP-265 the table attribute of `hive_storage_desc` `isOptional` is set to `true`. Post migration, `isOptional` set to `False`.

None.

CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command

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

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

In converting content from Navigator to Atlas, the create time for Hive tables is not moved to Atlas.

None

CDPD-10574: Suggestion order doesn't match search weights

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

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@58112
bc4 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) does not 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 the metadata collection from HMS. For example, if an ALTER TABLE operation in Spark that is changing a table name and is reporting 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 "@" does not fetch the correct results

When searching Atlas qualifiedName values that include an "at" 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, and use 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 does 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 logout 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 CDP.

None.

CDPD-1664: Guest users are redirected incorrectly

Authenticated users logging in to Atlas are redirected to the CDP 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: The 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, to add metadata tags to relationships that you wanted to make sure were unique in the system, you can design the relationship attribute with the property IsUnique equal true. However, in this release, the IsUnique attribute is not enforced.

None.

DOCS-13759: Tag Propagation stops after a certain depth while the lineage is being extended

When a tag is added to an entity at timestamp T1, the entities along the lineage to which the tag must be propagated is calculated at T1. Before tag propagation completes, if the lineage is extended, tag does not propagate to the entities in the extended lineage.

CDPD-41142: When a Kafka console consumer group is run, more than one update audits are seen

After running the console consumer with a consumer group, verify the consumer group entity created, along with the metrics and notifications for the consumer group and topic. The expected result can be: one ENTITY_CREATE audit and one ENTITY_UPDATE audit. But more than one ENTITY_UPDATE audits are seen.

CDPD-40165: Two audits are created for SPARK CTAS table

When following Spark queries are fired:

```
spark.sql("create table table1(id int)")
```

```
spark.sql("create table table2 as select * from table1")
```

HMS sends "ENTITY_CREATE" and "ENTITY_FULL_UPDATE_V2".

The extra ENTITY_FULL_UPDATE_V2 message received from HMS is sent as a part of ALTER_TABLE_ADDCOLS event from the HMS Hook side. This behaviour is observed only when the queries are run from Spark SQL and not when run the same queries are run from Beeline.

CDPD-39197: Debug metrics returns empty data

When debug metrics is enabled and some operations are performed, the response is empty

CDPD-36495: Updating legacyAttribute from False to True resets the initially created relationshipAttributes values

Creating types, entities, and to start, you must set the relationship with is_legacy_attribute value as **False**.

Later, update the value relationshipDef is_legacy_attribute to True.

For the entities that were created before updating the is_legacy_attribute to True, the relationship Attributes value is reset.

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-29409: Hive import: Suggestion suggests entity which is deleted.

Suggestions suggests tables of a database, which is a deleted entity.

CDPD-25152: Tag propagation through deferred actions consumes additional time as compared to default flow

The additional time might be due to the small overhead added to create / update task vertex and which is run in the background. This also depends on number of tasks queued to be executed.

CDPD-42954: Zeppelin notebook fails after enabling Atlas-HDFS hook

The Zeppelin notebooks are failing with errors after enabling Atlas-HDFS hook in the CDP cluster.

When the below properties are set for atlas-client.properties in Cloudera Manager:

- atlas.jaas.KafkaClient.option.keyTab
- atlas.jaas.KafkaClient.option.principal

Along with adding the properties in /etc/atlas/conf/atlas-application.properties, Cloudera Manager also adds these properties to atlas-application.properties for other services (like Spark).

Adding these properties interferes with the normal flow of the services (like Spark)

To enable HDFS lineage feature, instead of setting these properties through Cloudera Manager, users can manually add the properties directly in /etc/atlas/conf/atlas-application.properties

CDPD-10576: Deleted Business Metadata attributes appear in Search Suggestions

Atlas search suggestions continue to show Business Metadata attributes even if the attributes have been deleted.

None.

CDPD-39427: [HDFS Lineage]: When the input is a directory in case of put/copyfromLocal/cp/mv commands, lineage is not created even though the script succeeds.

When Source is a directory and target is a directory which is already present in Atlas, the command succeeds and inserts the data in the desired location, but lineage is not created.

DOCS-13760: System Attributes search, __classificationNames: Search with parent tag does not return entities associated to its children tags

System attribute search with __classificationNames = parent_tag returns entities associated to parent_tag only and not entities associated to its children tag.

Workaround: Instead of using system attribute, employ the basic search attribute "classification" which lists entities associated with inherited classifications.

CDPD-35818: Basic search with tag filter provides approximateCount as -1 when there is no match and is 0 otherwise

When the following search operations are performed:

- Faceted search with both tag filter and entity filter

The observed approximateCount is -1.

Here both entity filter and tag filter are present and when there is no match the response received is -1.

- Faceted search with only entity filter

Performing a basic search provides an approximate value of 0 when there is no match.

- Faceted search with only tag filter

Whenever there is tag filter in the query and there is no entity match, the approximateCount is -1 and if the tag filter is not available, the response approximateCount is 0

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export

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

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](#)

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

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

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

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

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

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

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

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

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-88230: Bulk import operations currently do not construct lineage as expected

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

Bulk import operations do not create audits for the imported entities.

Known Issues in Cruise Control

Learn about the known issues in Cruise Control, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

Rebalancing with Cruise Control does not work due to the metric reporter failing to report the CPU usage metric

On the Kafka broker, the Cruise control metric reporter plugin may fail to report the CPU usage metric.

If the CPU usage metric is not reported, the numValidWindows in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.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 the [Moving a Host Between Clusters](#).



Note: Cluster nodes affected by this issue are not displayed as unhealthy.

Known Issues in Hive

Learn about the known issues in Hive, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

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

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

The following services support DDL statements:

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

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

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

CDPD-60770: Beeline Authentication Issue with Special Characters in Passwords

When LDAP is enabled, users cannot authenticate with Beeline if the password contains a special character. For example, the following string fails:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>;user=user@XXX;password='R3G#xpXyoy1M0Jb1'
```

Use the `-p` parameter to execute the Beeline command:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>; -n user@XXX -p 'R3G#xpXyoy1M0Jb1'
```

CDPD-43769: Inconsistent behavior with CHAR comparisons and string literals

In Hive, 'char' data to 'string' data comparison is considered as equal, but 'char' data to 'varchar' data comparison is considered as not equal if the data has trailing spaces. For example,

```
create table test(colchar char(10), colvarchar varchar(20), colstring string);
insert into test values ('a ', 'a ', 'a ');

select count(*) from test where colchar = colstring;
Output returns 1 because char and string data is considered as equal

select count(*) from test where colchar = colvarchar;
Output returns 0 because char and varchar data is considered as not equal
```

It is recommended that you use 'varchar or string' data type instead of 'char' data type for any character or string representation.

CDPD-42726: Wrong results for agg queries from stats after running multi-insert query on managed table

Multiple inserts at the same time into the same partition results in invalid stats. For example:

```
create table source(p int, key int, value string);
insert into source(p, key, value) values (101,42,'string42');

create table stats_part(key int,value string) partitioned by (p int);
from source
insert into stats_part select key, value, p
insert into stats_part select key, value, p;

select count(*) from stats_part;
-- In this scenario, StatsOptimizer helps serving this query because the result should be rowNum of the partition p=101.
-- The result is 1, however, the result should be 2.
```

Rewrite the multi-insert statement by merging the branches inserting into the same partition to one branch.

CDPD-57574: Query execution fails due to NullPointerException in DagUtils.setupQuickStart

Hive queries may fail at runtime with a `NullPointerException` while executing the Tez graph at `DagUtils.setupQuickStart` method. A part of the stack trace from the error is shown below:

```
ERROR : Failed to execute tez graph.
java.lang.NullPointerException: null at org.apache.hadoop.hive.ql.exec.tez.DagUtils.setupQuickStart(DagUtils.java:1724)
```

The error usually occurs from an invalid plan that has cycles created by semi-join or map-join edges.

Disable some of the optimizations in order to remove the cycles from the plan. It is sufficient to set one of the following properties:

- `set hive.tez.dynamic.partition.pruning=false;`
- `set hive.tez.dynamic.semijoin.reduction=false;`
- `set hive.auto.convert.join=false;`

To avoid any undesired performance regressions, disable the optimization(s) only for specific queries and not globally (through Cloudera Manager) for the whole workload or cluster.

CDPD-45370: Cartesian join fails with an empty table

When the configuration `hive.tez.cartesian-product.enabled` is set to true, query fails if one of the tables involved in a Cartesian join has zero rows. This results in an `IllegalArgumentException` with the message `Number of items is 0. Should be positive.`

To address this issue, set `hive.tez.cartesian-product.enabled` to false before running the query.

Apache Jira: [\[HIVE-26582\]](#)

CDPD-43107: SemanticException for INSERT INTO statement when Hive Cost-based Optimizer (CBO) is disabled

If you are running the `INSERT INTO` query and `hive.cbo.enable` is set to "false", the query fails with a `SemanticException`. For example,

```
set hive.cbo.enable=false;

CREATE TABLE mytable (
  id INT,
  str STRING
);

INSERT INTO mytable (id, str) VALUES (1, 'a');
```

```
Output:
org.apache.hadoop.hive.ql.parse.SemanticException: 0:0 Expected
2 columns for insclause-0/default@mytable; select produces 1 col
umns. Error encountered near token 'a'
```

This issue occurs only when columns are specified for the target table in the query. `INSERT INTO mytable values (1, 'a');` does not result in an exception.

Enable CBO and run the query.

CDPD-43957: HiveServer shuts down during replication due to high resource usage

During Hive replication (both bootstrap and incremental), you may notice that the HiveServer (HS2) shuts down periodically with the following error:

```
java.sql.SQLException: org.apache.hive.jdbc.ZooKeeperHiveClientE
xception: Unable to read HiveServer2 configs from ZooKeeper
at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.ja
va:265)
at org.apache.hive.jdbc.HiveDriver.connect(HiveDriver.java:107)
```

```

at java.sql.DriverManager.getConnection(DriverManager.java:664)
at java.sql.DriverManager.getConnection(DriverManager.java:247)
at com.cloudera.enterprise.hive3qt.Hive3QueryTool$HiveOperation.
execute(Hive3QueryTool.java:682)
at com.cloudera.enterprise.hive3qt.Hive3QueryTool.main(Hive3Qu
eryTool.java:935)
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException:
Unable to read HiveServer2 configs from ZooKeeper
at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureCo
nnParams(ZooKeeperHiveClientHelper.java:177)
at org.apache.hive.jdbc.Utills.configureConnParamsFromZooKeeper
(Utills.java:580)
at org.apache.hive.jdbc.Utills.parseURL(Utills.java:391)
at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.j
ava:263)
... 5 more
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: T
ried all existing HiveServer2 uris from ZooKeeper.
at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.getServerHosts
(ZooKeeperHiveClientHelper.java:132)
at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureCon
nParams(ZooKeeperHiveClientHelper.java:172)
... 8 mor

```

This issue occurs when you replicate at scale with an unbalanced cluster setup that has all the roles running on the Cloudera Manager host. As a result, Cloudera Manager ends up in a bottleneck situation because HS2 crashes, preventing further replication. The logs (/var/log/messages) indicate that there were a large number of Java processes running on the host, which exhausted the host's memory and triggered an Out Of Memory Killer process to stop HS2.

You can restart HS2 and reinitiate replication. However, the replication process may take a longer time to complete.

- Follow these recommendations when setting up your source and target clusters:
 - Ensure that no HS2, Hive metastore (HMS), DataNode, or NameNode roles are running on the Cloudera Manager host.
 - Ensure that you have multiple instances of HS2 and HMS roles on different nodes.
 - It is recommended that you have the HS2 and HMS roles on different nodes than the DataNodes or NameNodes.

These practices increase the possibility of Hive replication at scale completing successfully.

- Limit the scale of Hive replication by temporarily disabling replication policies.

CDPD-40730: Parquet change can cause incompatibility

Parquet files written by the parquet-mr library in CDP 7.1.9, where the schema contains a timestamp with no UTC conversion will not be compatible with older versions of Parquet readers. The effect is that the older versions will still consider these timestamps as they would require UTC conversions and will thus end up with a wrong result. You can encounter this problem only when you write Parquet-based tables using Hive, and tables have the non-default configuration `hive.parquet.write.int64.timestamp=true`.

None.

CDPD-41274: HWC + Oozie issue: Could not open client transport with JDBC Uri

Currently only Spark cluster mode is supported in the Oozie Spark Action with Hive Warehouse Connector (HWC).

Use Spark action in cluster mode.

```

<spark xmlns="uri:oozie:spark-action:1.0">
  ...

```

```
<mode>cluster</mode>
...
</spark>
```

CDPD-26556 After an upgrade, querying a CTAS table under certain conditions might throw an exception

If you upgrade your Hive cluster from CDH 6 to CDP 7, create a CTAS table in the CDP cluster from a table you upgraded from CDH, you might see the following exception when you query the new table:

```
class org.apache.hadoop.io.IntWritable cannot be cast to class o
rg.apache.hadoop.hive.serde2.objectinspector.StandardUnionObject
Inspector$StandardUnion
```

This issue involves CDH-based tables having columns of complex types ARRAY, MAP, and STRUCT.

CDPD-23506: OutOfMemoryError in LLAP

Long running spark-shell applications can leave sessions in interactive Hiveserver2 until the Spark application finishes (user exists from spark-shell), causing memory pressure in case of a high number of queries in the same shell (1000+).

You must close spark-shell so that sessions are closed. Add the owner of the database or the tables as a user with read or read/write access to the tables directly.

CDPD-23041: DROP TABLE on a table having an index does not work

If you migrate a Hive table to CDP having an index, DROP TABLE does not drop the table. Hive no longer supports indexes ([HIVE-18448](#)). A foreign key constraint on the indexed table prevents dropping the table. Attempting to drop such a table results in the following error:

```
java.sql.BatchUpdateException: Cannot delete or update a parent
row: a foreign key constraint fails ("hive"."IDXS", CONSTRAINT "
IDXS_FK1" FOREIGN KEY ("ORIG_TBL_ID") REFERENCES "TBLS" ("TBL_ID"
))
```

There are two workarounds:

- Drop the foreign key "IDXS_FK1" on the "IDXS" table within the metastore. You can also manually drop indexes, but do not cascade any drops because the IDXS table includes references to "TBLS".
- Launch an older version of Hive, such as Hive 2.3 that includes IDXS in the DDL, and then drop the indexes as described in [Language Manual Indexing](#).

Apache Issue: [Hive-24815](#)

CDPD-17766: Queries fail when using spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive.

Do not specify spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive remotely.

Workaround: specify principal=hive.server2.authentication.kerberos.principal as shown in the following syntax:

```
jdbc:hive://<host>:<port>/<dbName>;principal=hive.server2.authen
tication.kerberos.principal;<otherSessionConfs>?<hiveConfs>#<hiv
eVars>
```

CDPD-13636: Hive job fails with OutOfMemory exception in the Azure DE cluster

Set the parameter hive.optimize.sort.dynamic.partition.threshold=0. Add this parameter in Cloudera Manager (Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml)

CDPD-10848: HiveServer Web UI displays incorrect data

If you enabled auto-TLS for TLS encryption, the HiveServer2 Web UI does not display the correct data in the following tables: Active Sessions, Open Queries, Last Max n Closed Queries

Known Issues in HDFS

Learn about the known issues in HDFS, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

There are no known issues identified in the releases before Cloudera Runtime 7.1.9 SP1 CHF 12.

Known Issues in Hue

Learn about the known issues in Hue, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12**CDPD-58978: Batch query execution using Hue fails with Kerberos error**

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark queries in a batch mode launches a shell script Oozie job from Hue and this is not supported on a secure cluster.

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

OPSAPS-69659: Hue service fails on restart with "Unable to find pycogp2 2.5.4" error

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.

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

Key: PYTHONPATH

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

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

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

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

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

None.

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

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

None.

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

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

None.

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

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

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

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

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

CDPD-43293: Unable to import Impala table using Importer

Creating Impala tables using the Hue Importer may fail.

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

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

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

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

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

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

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

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

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

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

Hue 4 web interface uses the unsafe-inline directive in its CSP header. As a result, the application server does not set the CSP header in its HTTP responses, and therefore does not benefit from the additional protection against potential cross-site scripting issues and other modern application vulnerabilities which a properly configured CSP may provide. This could lead to application vulnerability.

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

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

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

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

```
[search]
# URL of the Solr Server
solr_url=http://[***HOSTNAME***]:[***PORT***]/solr/
```

For example:

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

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

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

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

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

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

Invalid S3 URI error while accessing S3 bucket

The Hue Load Balancer merges the double slashes (//) in the S3 URI into a single slash (/) so that the URI prefix "/filebrowser/view=S3A://" is changed to "/filebrowser/view=S3A:". This results in an error when you try to access the S3 buckets from the Hue File Browser through the port 8889.

The Hue web UI displays the following error: “Unknown error occurred”.

The Hue server logs record the “ValueError: Invalid S3 URI: S3A” error.

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.

You should be able to load the S3 browser from both 8888 and 8889 ports.

Alternatively, you can use the Hue server port 8888 instead of the load balancer port 8889 to resolve this issue.

Error while rerunning Oozie workflow

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

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

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

CDPD-17956, CDPD-14167: Setting idle session timeout for Hue does not work when the cluster is secured using Knox SSO

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

None

Query history displays "Result Expired" message

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

None.

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

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

Hue limitation after upgrading from CDH to CDP Private Cloud Base

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

Hue UI is blank upon login after upgrading to CDP 7.1.7 from CDH 6

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

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

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



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

4. Specify the Knox Gateway hostname in the Knox Proxy Hosts field.

5. Click Save Changes and restart the Hue service.

Unsupported features

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

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

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

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

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

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

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

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

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

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



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

Known Issues in Impala

Learn about the known issues in Impala, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

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

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

The following services support DDL statements:

- Impala

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

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

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

CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the `--tls_ciphersuites` startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.

Ensure that the list in the `--tls_ciphersuites` startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

Apache Jira: [IMPALA-14625](#)

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work

If you try to run `impala-shell` on RHEL 9 by setting the default python executable available in `PATH` to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, `IMPALA_PYTHON_EXECUTABLE=python3`.

CDPD-81980: Impala query may fail due to backpressure

Impala queries could fail with a dropped due to backpressure error, even when service queue and memory usage appear normal.

Increase the value of the `--mem_limit` startup option.

CDPD-42958: After upgrading the CDH 7.1.9 from CDH 6.x, under certain conditions you cannot insert data into a table

Under the following conditions, after upgrading from CDH 6.x to CDH 7.1.9 you cannot insert data into a table from Impala:

- On CDH 6.x, you created a database with Impala in a user specified HDFS location.
- Using Hive, you then created a table in the database.

Under these conditions, the database and table are stored in the user-specified HDFS directory. After upgrading, the HDFS directory of the table is read-only for Impala. Consequently, from Impala you cannot insert new data into the table because Impala does not have write permission on the HDFS directory.

Workaround: To resolve this issue, use either one of the following workarounds:

- Using the Ranger Web UI, in the policy repository `cm_hdfs`, grant the user 'impala' write permission on the directory where the table resides.
- Enter the following command to grant write permission to user 'impala' on the HDFS directory where the table resides.

```
hdfs dfs -setfacl -m default:user:impala:rwx <HDFS directory>
```

CDPD-80019: Impala fails to start due to large catalog-update topic

Impala coordinators fail to start if the catalog-update statestore topic exceeds 2GB, hitting the Thrift 0.16.0 message size limit (`thrift_rpc_max_message_size`). This prevents the coordinator from processing the update.

- Enable local catalog mode: `--use_local_catalog=true`
- In Cloudera Manager, go to Impala Configuration search for "Catalog Server Command Line Argument Advanced Configuration Snippet (Safety Valve)" and click 
- Add property name and value:

```
Name : --catalog_topic_mode
Value: minimal
```

- Save changes.

Apache Jira: [IMPALA-12350](#)

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 HDF
S file hdfs://nameservice1/warehouse/tablespace/managed/hive/<da
tabase>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException:
File does not exist: /warehouse/tablespace/managed/hive/<data
base>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

You must use a single Knox node to access Impala UI.

Impala api calls via kinox require configuration if the kinox customized kerberos principal name is a default service user name

To access impala api calls via kinox, if the kinox customized kerberos principal name is a default service user name, then configure "authorized_proxy_user_config" by clicking Clusters->impala->configuration. Include the kinox customized kerberos principal name in the comma separated list of values `<knox_custom_kerberos_principal_name>=*` where `<knox_custom_kerberos_principal_name>` is the value of the Kerberos Principal in the Knox service. Select Clusters>Knox>Configuration and search for Kerberos Principal to display this value.

CDPD-21828: Multiple permission assignment through grant is not working

None

Problem configuring masking on tables using Ranger

The following Knowledge Base article describes the behavior when we configure masking on tables using Ranger. This configuration works for Hive, but breaks queries in some scenarios for Impala.

For a workaround, see the following Knowledge Base article: [ERROR: "AnalysisException: No matching function with signature: mask\(FLOAT\)" when Impala jobs fail with the following error with signature: mask\(FLOAT\)](#)

IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXXXX"

IMPALA-532: Impala should tolerate bad locale settings

If the LC_* environment variables specify an unsupported locale, Impala does not start.

Add LC_ALL="C" to the environment settings for both the Impala daemon and the Statestore daemon.

Avro Scanner fails to parse some schemas

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

IMPALA-691: Process mem limit does not account for the JVM's memory usage

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the impalad daemon.

To monitor overall memory usage, use the top command, or add the memory figures in the Impala web UI /memz tab to JVM memory usage shown on the /metrics tab.

IMPALA-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Remove trailing semicolon from the Avro schema.

IMPALA-1652: Incorrect results with basic predicate on CHAR typed column

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the RPAD() function to blank-pad literals compared with CHAR columns to the expected length.

IMPALA-1821: Casting scenarios with invalid/inconsistent results

Using a CAST() function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

IMPALA-2005: A failed CTAS does not drop the table if the insert fails

If a CREATE TABLE AS SELECT operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

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

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.

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.

Known Issues in Iceberg

Learn about the known issues in Iceberg, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

There are no known issues identified in earlier releases.

Known Issues in Kafka

Learn about the known issues in Kafka, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

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

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

The following services support DDL statements:

- Impala

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

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

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

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker) in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Save Changes > Restart SMM.

The offsets.topic.replication.factor property must be less than or equal to the number of live brokers

The offsets.topic.replication.factor broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a GROUP_COORDINATOR_NOT_AVAILABLE error until the cluster size meets this replication factor requirement.

None

Requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true

The first few produce requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true.

Increase the number of retries in the producer configuration setting retries.

Performance degradation when SSL Is enabled

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with ssl.secure.random.implementation = SHA1PRNG. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

OPSAPS-43236: Kafka garbage collection logs are written to the process directory

By default Kafka garbage collection logs are written to the agent process directory. Changing the default path for these log files is currently unsupported.

None

RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transactional method because we are in an error state
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeFailWithError(TransactionManager.java:1125)
```

```

        at org.apache.kafka.clients.producer.internals.TransactionManager.maybeAddPartition(TransactionManager.java:442)
        at org.apache.kafka.clients.producer.KafkaProducer.doSend(KafkaProducer.java:1000)
        at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:914)
        at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:800)
        .
        .
        .
    Caused by: org.apache.kafka.common.errors.ClusterAuthorizationException: Cluster authorization failed.

```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting `enable.idempotence` to `false`.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

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

The Debezium Db2 Source connector does not support the evolution (updates) of schemas. In addition, schema change events are not emitted to the schema change topic if there is a change in the schema of a table that is in capture mode. For more information, see [DBZ-4990](#).

None.

CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy compression. The HDFS Stateless Sink connector is only affected if the `Compression Codec` or `Compression Codec` for Parquet properties are set to `SNAPPY`.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: org.apache.hadoop.util.NativeCodeLoader.buildSupportsSnappy()
```

```
Failed to write to HDFS due to java.lang.RuntimeException: native snappy library not available: this version of libhadoop was built without snappy support.
```

Download and deploy missing libraries.



Important: Ensure that you complete steps 1-11 on all Kafka Connect hosts. Additionally, ensure that the advanced configuration snippet in step 12 is configured for all Kafka Connect role instances.

1. Create the `/opt/nativelibs` directory.

```
mkdir /opt/nativelibs
```

2. Change the owner to kafka.

```
chown kafka:kafka /opt/nativelibs
```

3. Locate the directory containing the Hadoop native libraries and copy its contents to the directory you created.

```
cp /opt/cloudera/parcels/CDH/lib/hadoop/lib/native/* /opt/nativelibs
```

4. Verify that 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

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect

instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Unsupported Features

The following Kafka features are not supported in Cloudera Data Platform:

- Only Java and .Net based clients are supported. Clients developed with C, C++, Python, and other languages are currently not supported.
- The Kafka default authorizer is not supported. This includes setting ACLs and all related APIs, broker functionality, and command-line tools.
- SASL/SCRAM is only supported for delegation token based authentication. It is not supported as a standalone authentication mechanism.
- Kafka KRaft in this release of Cloudera Runtime is in technical preview and does not support the following:
 - Deployments with multiple log directories. This includes deployments that use JBOD for storage.
 - Delegation token based authentication.
 - Migrating an already running Kafka service from ZooKeeper to KRaft.
 - Atlas Integration.

Limitations

Collection of Partition Level Metrics May Cause Cloudera Manager's Performance to Degrade

If the Kafka service operates with a large number of partitions, collection of partition level metrics may cause Cloudera Manager's performance to degrade.

If you are observing performance degradation and your cluster is operating with a high number of partitions, you can choose to disable the collection of partition level metrics.



Important: If you are using SMM to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both SMM and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, SMM will not be able to display information about partitions. In addition, Cruise Control will not operate properly.

Complete the following steps to turn off the collection of partition level metrics:

1. Obtain the Kafka service name:
 - a. In Cloudera Manager, Select the Kafka service.
 - b. Select any available chart, and select Open in Chart Builder from the configuration icon drop-down.
 - c. Find \$SERVICENAME= near the top of the display.
The Kafka service name is the value of \$SERVICENAME.
2. Turn off the collection of partition level metrics:
 - a. Go to Hosts Hosts Configuration .
 - b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

Enter the following to turn off the collection of partition level metrics:

```
[KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_entity_update_enabled=false
```

Replace [KAFKA_SERVICE_NAME] with the service name of Kafka obtained in step 1. The service name should always be in lower case.

- c. Click Save Changes.

Known Issues in Knox

Learn about the known issues in Knox, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

CDPD-76435: Default HTTP client timeout for Streams Messaging Manager requests may cause connection timeouts

Some Streams Messaging Manager requests, particularly those retrieving topic data, may take up to 30 seconds to complete, even when there are no connectivity issues. Since the default gateway timeout is set to 20 seconds, this can result in connection timeouts despite the system functioning as expected.

Increase the HTTP client and socket timeouts for Streams Messaging Manager services to 60 seconds. This prevents connection timeouts when the requests take longer than the default gateway timeout.

Follow these steps to configure the required properties for the SMM API and SMM UI services:

1. In Cloudera Manager, select the Knox service.
2. Go to Configuration.
3. Find the Knox Simplified Topology Management - cdp-proxy property, and click **+** to add the following lines:

```
SMM-UI:replayBufferSize=512
SMM-UI:httpclient.connectionTimeout=1m
SMM-UI:client.socketTimeout=1m
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

4. Find the Knox Simplified Topology Management - cdp-proxy-api property, and click **+** to add the following lines:

```
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

5. Find the Knox Simplified Topology Management - cdp-proxy-token property, and click **+** to add the following lines:

```
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

6. Save your changes.

7. Refresh the cluster.

Known Issues in Kudu

Learn about the known issues in Kudu, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

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.

Workaround: Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

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.

Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

Unsupported features

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 in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

The following known issues were identified in this release:

No new known issues in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

No known issues identified in Livy.

Known issues in Navigator Encrypt

Learn about the known issues in Navigator Encrypt, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

Failed to start navencrypt-mount.service

After installing NavEncrypt on SLES, if the command "systemctl status navencrypt-mount" fails with the error: "Failed to start navencrypt-mount.service: Unit navencrypt-mount.service failed to load: No such file or directory", the systemd files need to be reloaded.

Run the command:

```
(sudo) systemctl daemon-reload
```

Crash in Navigator Encrypt kernel module if `--profile-file=` option issued with ACLs

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 Oozie

Learn about the known issues in Oozie, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

There are no known issues identified in the releases before Cloudera Runtime 7.1.9 SP1 CHF 12.

Known Issues in Ozone

Learn about the known issues in Ozone, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

CDPD-87429: Unexpected log lines in Ozone snapshot diff output

``ozone.client.server-defaults.validity.period.ms`` configuration introduced to get server default configurations causes Replication Manager to fail when replicating data between Ozone buckets with the following error when the config value does not specify a timeunit.

```

` ``
java.lang.IllegalStateException: Invalid first line in diff: "25
/07/02 05:58:44 INFO Configuration.deprecation: No unit for ozon
e.client.server-defaults.validity.period.ms(3600000) assuming MI
LLISECONDS"
  at com.cloudera.enterprise.distcp.OzoneCopyListingUtils.parseOz
oneDiff(OzoneCopyListingUtils.java:397)
` ``

```

Overwrite the ``ozone.client.server-defaults.validity.period.ms`` to configure the value to have time unit (ns, ms, s, m, h). For example, 3600000ms.

CDPD-71648: Snapshot creation failing with "PERMISSION_DENIED" error on a bucket created during pre-upgrade in 7.1.7 SP3

On upgrade of cluster from 7.1.7 SP3 to 7.1.9 SP1, we are not able to create snapshots on a pre-upgrade volume or bucket with "PERMISSION_DENIED" error, since in 7.1.7SP3, we were using full KRB owner name with domain but currently, we only use and search for user shortname.

Update the owner of the bucket with the user shortname and then create snapshot by running the commands:

```
ozone sh bucket update --user=[***USER_SHORT_NAME***] {volume}/
{bucket}
```



```
ozone sh snapshot create {volume}/{bucket}
{[***SNAPSHOT_NAME***]}
```

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

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

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

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

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

CDPD-64865: Intermittent timeout in TestBlockDeletion.testBlockDeletion (BlockDeletionService stuck)

If the event of SCM exiting safe mode was triggered multiple times, this might eventually cause the SCMBlockDeletingService to transition into the PAUSING state.

Check the SCM logs and ensure that the SCMBlockDeletingService is not running, and then restart the SCM.

CDPD-81569: Recon is not reflecting correct metrics for Volumes, Buckets, Keys, and Pending Key Deletions

Recon is displaying data for container and space utilisation only. The metrics for Volumes, Buckets, Keys, and Pending Key Deletions are not being updated and displayed correctly.

Apache Jira: [HDDS-12607](#)

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: 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.

The Prometheus binary is not available in CDP Private Cloud Base 7.1.9 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.

After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and if Ozone is in the Non-HA environment, an exception message is observed during the finalization of the Ozone upgrade.

During the finalization of the upgrade, a ClassNotFoundException message for org.cloudera.log4j.redactor.RedactorAppender class may be displayed. The error message can be ignored because the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

CDPD-68951: In CDP Private Cloud Base 7.1.9 CHF2 version and lower, the command `ozone sh key list <bucket_path>` displays the isFile flag in a key's metadata as false even when the key is a file. This issue is rectified in CDP Private Cloud Base 7.1.9 CHF3. However, the pre-existing (pre-upgrade) key's metadata cannot be changed.

None

CDPD-67771: When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(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.

CDPD-64398: In SCM with non-HA configuration, Secret key manager does not initialise. Hence, the startup of OM and DataNode fails because it cannot get a secret key. This key is used when security is enabled and used in Block token and container token verification while communicating between the Ozone client, OM, and Datanode.

This issue does not occur in SCM with HA configuration.

You must force exit from the safe mode for SCM. This triggers the initialization of the Secret key manager. Below are the options:

- exit safe mode manually whenever SCM is started
- disable safe mode by setting `hdds.scm.safemode.enabled=false` in safety valve for SCM configuration

Impact of disabling safe mode and the purpose of safe mode:

- Write should not fail once the cluster is out of safe mode
- Read of existing data should not fail after the cluster is out of safe mode
- Unnecessary re-replication should be avoided during cluster restart

So disabling safe mode does not have any major impact on Ozone.

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.

CDPD-52412: keyManagerImpl#listStatus exceeds the maximum RPC length.

To reduce the payload, you must increase the part size.

OPSAPS-68159: If you did not deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Error when distributing parcel to host, Ignoring non-compliant parcel manifest error is displayed after Cloudera Manager is upgraded to 7.11.3.

If you encounter the error, perform the following steps:

1. Deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.11.3.
2. Restart the cluster with a delay of 10 minutes.
3. Continue to upgrade CDH 7.1.8 to CDH 7.1.9.

OPSAPS-68159: If you did not deactivate the Ozone parcel 718.2.x on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Ozone roles during the CDH 7.1.8 upgrade to CDH 7.1.9.

If you encounter the error, perform the following steps:

1. Deactivate the Ozone parcel 718.2.x.
2. Restart the Ozone service.
3. Perform Finalize Upgrade for the Ozone service.

Step result: The Ozone roles will be displayed in green.

CDPD-60989: The packaging version for Apache Ozone points to the 1.2.0 older version. This is a version string problem and not a packaging issue. The version of the Apache Ozone binary is closest to 1.4.0.

None. This only affects the JAR versioning.

CDPD-60366: Native library loader fails when system property `native.lib.tmp.dir` is not set. It fails because the library is copied to / instead of the `cwd`.

Setting `native.lib.tmp.dir` to a certain path like `/tmp` should solve the issue. To set this system property, you must add `-Dnative.lib.tmp.dir=/tmp` to `ozone_java_opts` in Cloudera Manager configuration of the Ozone cluster.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-60598: DataNode can consume up to the JVM heap size worth of direct memory buffers under specific failure scenarios. This can lead to a DataNode crash. This is due to Netty allocating and not freeing the direct memory buffers.

Restarting DataNode will resolve this issue.

CDPD-60466: The Ozone client is missing in the Cloudera Manager's Service Monitor node. This causes Ozone Canary's health check to fail.

Ensure you install Ozone CLI on the same node as Cloudera Manager's Service Monitor.

CDPD-60012: The Ozone SCM may be stuck in safe mode after upgrade to 7.1.9. This is due to a bug in how the SCM accounts for open storage containers.

The workaround is to temporarily set the Ozone configuration `hdds.scm.safemode.threshold.pct` to a lower value like 0.90 and restart the SCM.

CDPD-59679: In some scenarios, excess UNHEALTHY replicas of EC containers may not be removed from the cluster.

None

CDPD-50116: Topology-aware reads can provide better performance by routing applications to the closest replica of a file, if available, on the same physical node or same rack. However this feature is disabled by default in CDP 7.1.9.

This feature can be enabled by setting `ozone.network.topology.aware.read` to true using Cloudera Manager safety valve.

CDPD-57165: The DataNode disk usage thread may abort due to Ratis log tailing.

Use `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory` to calculate disk usage. Configure in the `ozone-site.xml` property: `hdds.datanode.du.factory` value: `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory`

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.

CDPD-55513: Hive external table replication policy does not migrate Hive tables in HDFS storage to Ozone.

To replicate data from HDFS to Ozone, see [Migrating your data from HDFS to Ozone](#). You can use HMS Mirror to replicate the metadata.

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"
```

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_SHAREDLIB_DIR=/VAR/TMP`
3. Restart Ozone.

CDPD-49137: Ozone Manager Kerberos token expires for SCM communication and OM does not log in again.

Sometimes, OM's Kerberos token is not updated and it stops to communicate with SCM. When this occurs, writes start failing.

Restart OM or set the safety valve `hadoop.kerberos.keytab.login.autorenewal.enabled = true`

CDPD-56684: Keys get deleted when you do not have permission on volume

When a volume is deleted, it recursively deletes the buckets and keys inside it and only then deletes the volume. The volume delete ACL check is done only in the end, due to which you may end up deleting all the data inside the volume without having delete permission on the volume.



Note: There was no bucket/key permission available which allowed the user to delete them recursively.

CDPD-50610: Large file uploads are slow with OPEN and stream data approach

Hue file browser uses the append operation for large files. This API is not supported by Ozone in 7.1.9 and therefore large file uploads can be slow or timeout from the browser.

Use native Ozone client to upload large files instead of the Hue file browser.

OPSAPS-64097: Ozone service restart failed at SCM

Stopping SCM service using Cloudera Manager can sometimes timeout and need a retry. The Cloudera Manager API waits for 90 seconds which is not sufficient under certain circumstances.

Retry the shutdown using Cloudera Manager if the SCM still shows as running after a refresh of the service status.

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-filesystem-hadoop3.jar`, it is likely due to the presence of the user `hdfs` on the node prior to CDP parcel activation. As a result, the Cloudera Manager agent skips the initialization related to HDFS, leading to the non-creation of the `/var/lib/hadoop-hdfs` directory.

Create the directory `“/var/lib/hadoop-hdfs”` on all nodes followed by the deactivation and activation of the CDP parcel (deactivate and activate the Ozone parcel instead, in case Ozone parcel is used).

CDPD-50447: When SCM High Availability is enabled, each of the SCM web UIs report the host of the web UI as the leader of HA, and the other two as followers. This gives wrong information

Correct output is available by running the `ozone admin scm roles --service-id=<ID>` command.

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](#).

OPSAPS-62327: In an Ozone cluster without any gateway roles, Ozone is unable to deploy client configurations and displays the ConfigGenException error.

You must add the Ozone gateway roles to the cluster.

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-36389: The configurations `datanodes.involved.max.percentage.per.iteration` and `size.moved.max.per.iteration` are meant to limit the max number of DataNodes that is involved and maximum size that can move in an iteration. This bug causes balancer to stop an iteration when it is two DataNodes or 1 Container size (5GB) away from hitting these limits. However, these DataNodes can again be considered for balancing in the next iteration. This means the cluster is balanced after enough iterations, albeit a bit slowly. This bug is apparent in small clusters of around four DataNodes where the DataNodes can be either the source or target for a lot of moves but the iteration gets stopped when three DataNodes have been involved. It takes a higher number of iterations to eventually balance this cluster. While this is a performance issue, it does not prevent balancer from ultimately balancing the cluster.

Increase the speed for balancing by decreasing the interval between each iteration using the configuration `balancing.iteration.interval`. The value of this configuration must be greater than `hdds.datanode.du.refresh.period`. `size.moved.max.per.iteration` can be increased to allow more data to move in one iteration.

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-35141: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2> (state=08S01,code=40000) java.sql.SQLException: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2>. This issue does not occur if the source and target buckets are different in Hive queries. Currently, copying across the same bucket is only supported.

Avoid different buckets in source and target path. Ensure that `hive.exec.stagingdir` is unset or set to the same bucket as the warehouse so that staging directories are created in the same bucket as the warehouse.

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-41184: With Legacy buckets, FileSystem Optimised is not interoperating with the Ozone shell command. The cause is that the directory key entry in the DB KeyTable is stored as "dir1/" with trailing slash. But while performing the described operation, Ozone shell (o3://) is normalizing the given path and removing the trailing slash "/" from it. That results in KEY_NOT_FOUND exception.

There are three workarounds:

- Use FileSystem API to Delete the Directories rather than Shell-Command API.

- Use FSO buckets instead of Legacy Buckets. As in FSO, you can create Intermediate Directories and Delete Directories using the Ozone shell commands.
- Disable and set the configuration `ozone.om.enable.filesystem.pathsflag` to false in order to delete the directories. This is generally not a preferred workaround because the cluster must be restarted to pick up the new changes.

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-12542: Mount of Ozone filesystem with the help of FUSE fails.

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-42897: EC writes are failing with "No enough datanodes to choose" after EC replication config is set globally.

EC writes start failing when large number of pipelines are created as a result of multiple EC configs with different chunk sizes used to write keys.

If standard EC configs (that is, rs-3-2-1024k) are used to write keys, number of pipelines created per DataNode is limited to five and this issue is not seen with standard EC configs.

The recommendation is not to create too many random chunk sizes. It is configurable because, users can decide based on their workload. But do not to have separate chunk sizes for each file.

CDPD-41539: The "No such file or directory" imessage is displayed when EC file is read using older ofs client.

You must upgrade the client before trying to read the key: vol1/ecbuck1/1GB_ec".

CDPD-40560: Filesystem Operations using Hadoop S3a connector on a FILE_SYSTEM_OPTIMIZED bucket fails. org.apache.hadoop.ozone.om.exceptions.OMException: Unable to get file status: volume: s3v bucket: fso key: test/

Do not run Hadoop S3a commands on a FILE_SYSTEM_OPTIMIZED bucket. Use OBJECT_STORE bucket layouts.

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.

CDPD-43432: Ozone Service in fault state in DataNode - Long Running setup.

Upgrade RocksDB to the latest version.

OPSAPS-63999: In the newly installed cluster, the Finish upgrade option is clickable.

None

OPSAPS-64648: Failed to start Ozone node using Cloudera Manager if default log path /var/log/hadoop-ozone does not exist. If this path does not exist, any Ozone nodes(for example SCM or DataNode) restart will fail.

Run the following command `sudo -u hdfs mkdir -p /var/log/hadoop-ozone` or replace hdfs with the user Ozone roles that are running.

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.

Known Issues in Ranger

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

Access issues with users having similar name as admin or keyadmin users

Users with USER role that have a similar name as admin or keyadmin users are able to query those admin and keyadmin users.

None.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

UnsupportedClassVersionError

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.

CDPD-68806: The Revoke operation for users belonging to a group or role permission does not function as expected.

List command is listing all the tables even when the user permission is revoked. And also the command does not add any deny policy to Ranger for that specific user.

For example:

user_1 is the part of group_1

user_2 is the part of role_1

Create a Ranger policy for group_1 and role_1 to access all the resources in HBase.

Try to revoke the permission for user_1 and user_2 using HBase shell global revoke command:

```
revoke 'user_1'
```

```
revoke 'user_2'
```

But there is no deny policy added for that user. And access is still granted to user_1 and user_2.

This behavior is currently not supported in HBase shell. Must be handled manually using the Ranger policy change.

CDPD-68739: The revoke command does not work when using the HBase shell

While using the HBase shell, running the revoke command does not cancel the user permission. Users are able to perform actions even after running the revoke command.

For example:

When a user is granted permission to a particular table using the HBase grant command.

```
hbase:003:0> grant 'user_1', 'RWCA', 'some_abc_table'
```

And then execute the global revoke command for that user:

```
hbase:005:0> revoke 'user_1'
```

Permission for that user 'user_1' for the table 'some_abc_table' not got revoked.

None.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-58860: After upgrading CDP-7.1.8 to CDP-7.1.9, cdp-proxy token missing from Knox-Ranger policy.

As part of [OPSAPS-67480](#) in 719 default ranger policy is added from cdp-proxy-token topology, so that after a new installation of CDP-7.1.9, the Knox-Ranger policy includes cdp-proxy-token. However, upgrades do not add cdp-proxy-token to cm_knox policies automatically.

Manually add cdp-proxy-token to the Knox policy, using Ranger Admin Web UI.

1. Log in to Cloudera ManagerRangerRanger Admin Web UI, as a Ranger administrator.
2. On Ranger Admin Web UIService ManagerResourceKnox, click cm_knox.
3. In Knox Policies, open the CDP Proxy UI, API and Token policy.
4. In Knox Topology*, add cdp-proxy-token.
5. Click Save.
6. Restart Ranger.

CDPD-58704: hadoop roll key / key delete command shows operation failed error when one KMS host is down, even when operation succeeds

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered kms instances. if one KMS instance is stopped (not deleted), client gets 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.

```
hbase:001:0> revoke 'hrt_11'
ERROR: org.apache.hadoop.hbase.coprocessor.CoprocessorException: HTTP 400 Error: processSecureRevokeRequest processing failed
    at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocessor.preRevoke(RangerAuthorizationCoprocessor.java:1309)
    at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocessor.preRevoke(RangerAuthorizationCoprocessor.java:1128)
    at org.apache.hadoop.hbase.master.MasterCoprocessorHost$162.call(MasterCoprocessorHost.java:1857)
    at org.apache.hadoop.hbase.master.MasterCoprocessorHost$162.call(MasterCoprocessorHost.java:1854)
    at org.apache.hadoop.hbase.coprocessor.CoprocessorHost$ObserverOperationWithoutResult.callObserver(CoprocessorHost.java:558)
    at org.apache.hadoop.hbase.coprocessor.CoprocessorHost.executeOperation(CoprocessorHost.java:631)
    at org.apache.hadoop.hbase.master.MasterCoprocessorHost.preRevoke(MasterCoprocessorHost.java:1854)
    at org.apache.hadoop.hbase.master.MasterRpcServices.revoke(MasterRpcServices.java:2740)
    at org.apache.hadoop.hbase.shaded.protobuf.generated.MasterProtos$MasterService$2.callBlockingMethod(MasterProtos.java)
    at org.apache.hadoop.hbase.ipc.RpcServer.call(RpcServer.java:387)
    at org.apache.hadoop.hbase.ipc.CallRunner.run(CallRunner.java:139)
    at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:369)
    at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:349)
```

None

CDPD-56741: Improvement in log message when jwtauth not used

The following exception is printed at startup only and it is not cluttering logs:

```
2023-05-30 06:18:40,127 ERROR org.apache.ranger.rms.security.RMSJwtAuthFilter:
quasar-pibgzl-1.quasar-pibgzl.root.hwx.site-startStop-1]:
Failed to initialize Ranger RMS JWT Auth Filter.
java.lang.Exception: RangerJwtAuthHandler:
Mandatory configs ('jwks.provider-url' & 'jwt.public-key')
are missing, must provide atleast one.
    at org.apache.ranger.authz.handler.jwt.RangerJwtAuthHandler.initialize(RangerJwtAuthHandler.java:84)
    ~[ranger-authn-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
    at org.apache.ranger.rms.security.RMSJwtAuthFilter.initialize(RMSJwtAuthFilter.java:77)
    ~[ranger-rms-common-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
```

```

        at jdk.internal.reflect.NativeMethodAccessorImpl.invoke
0(Native Method) ~[?:?]
        at jdk.internal.reflect.NativeMethodAccessorImpl.invoke(N
ativeMethodAccessorImp

```

This exception is logged, since mandatory JWT auth filter configs are not provided (expected for some environments like y-cloud) Even though JWT auth filter is failed to initialised, it falls back to kerberos auth, hence no impact from auth perspective

NA

CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup

This is a warning log printed in catalina.out file when Ranger RMS server is initialised. Following exception is observed only in Oracle 19 setup. FileNotFoundException: /usr/share/java/oraclepki.jar

NA

CDPD-48975: Ranger KMS KTS to KMS DB migration : keys with the same name but different case are not migrated

KMS keys are not case sensitive

No workarounds. Such keys combination are very rare and migration doc was updated to check such keys before starting the migration.

CDPD-42598: Kafka policy creation allowed with incorrect permissions.

When creating a Kafka policy from the UI, the permissions "Idempotent write" and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource

None.

CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload:

```

{"resourceName": "classification", "userInput": "", "resources":  {""classification": []}}

```

expectation is to return all the classifications . But the response has only "classification" Happens similarly for entity-label , entity-business-metadata.

None.

CDPD-40734: User allowed to insert data into a hive table when there is a deny policy on a table column.

A user is allowed to enter data into a table even if there is a deny policy present on one of the table columns.

Test scenario details:

```

Policy setup :-
policy 1 :- all access policy for hrt_qa, hive and impala
users
resources - database - * , table - * , column - *
users : hrt_qa, hive, impala
access - all access allowed
policy 2 :- policy on test_1.table_1 for hrt_5
users : hrt_5
resources : database - test_1, table - table_1, column - *
access :- all access allowed
policy 3 :- deny policy on test_1.table_1.c0 for hrt_5
users : hrt_5
resources : database - test_1, table - table_1, column -
c0

```

```
access - all access denied
data setup :-
database - test_1
table - table_1(c0 int, c1 int)
```

The user is able to insert data into the table.

None.

Known Issues in Schema Registry

Learn about the known issues in Schema Registry, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

CDPD-40380: Authorization checking issue when Kerberos is disabled

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the `ranger.plugin.schema-registry.service.name` field.
3. Replace `GENERATED_RANGER_SERVICE_NAME` with the actual name of the service.
4. Restart the Schema Registry service.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

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

Schema Registry uses a schema fingerprinting mechanism to differentiate between schemas. However, fingerprinting does not take into consideration the schema attributes of the field type. As a result, if you have two schemas where the only difference is that one has type attributes defined and the other does not, they will be considered identical by Schema Registry. For example, the following schemas are considered identical:

```
#Schema V1
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"local_timestamp_micros_long","type":"long"}]}

#Schema V2
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"local_timestamp_micros_long","type":{"type":"long","logicalType":"local-timestamp-micros"}}]}
```

Notice that the only difference is that in the second schema, the `local_timestamp_micros_long` field has a logical type specified. In cases like this, the new version of the schema is not created, the initial version is used. This is true even if the data that is being produced has a new schema version. The ID of the first schema version is used and is put in the serialized record. The new schema version is not created.

This issue is common when using change data capture (CDC) connectors like the Debezium connectors. This is because CDC connectors create schemas with the logical type decimal based on the column type in the database schema. For example:

```
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"database_column","type":{"type":"bytes","logicalType":"decimal","precision":64,"scale":0}}]}
```

If the database schema changes (for example, the column type), it is possible that only scale changes, which is a schema attribute.

```
{ "type": "record", "name": "schema_name", "namespace": "ns", "fields":
  [ { "name": "database_column", "type": { "type": "bytes", "logicalType":
    "decimal", "precision": 64, "scale": 1 } } ] }
```

In this case, even though scale changed to 1, the first version of the schema is used where scale is 0. As a result, the data is consumed with the wrong scale.

Avoid using logical types or other attributes. Alternatively, ensure that there are no changes in the logical types or other attributes between schema versions.

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

Schema Registry does not start if the address specified in the Load Balancer Address Ranger property does not end with a trailing slash (/).

Set the value of the RANGER_REST_URL Schema Registry environment variable to an address that includes a trailing slash.

1. In Cloudera Manager, select the Schema Registry service.
2. Go to Configuration.
3. Find the Schema Registry Server Environment Advanced Configuration Snippet (Safety Valve) property and add the following:

```
Key: RANGER_REST_URL
Value: [***RANGER REST API URL***]
```

Replace [***RANGER REST API URL***] with an address that can be used by Schema Registry to access Ranger. Ensure that the address ends with a trailing slash. For example: <http://ranger-1.cloudera.com:6182/>

4. Restart the Schema Registry service.

OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

1. Access the Ranger Console (Ranger Admin web UI).
2. Click the cm_atlas resource-based service.
3. Add the schemaregistry user to the all - * policies.
4. Click Manage Service Edit Service .
5. Add the schemaregistry user to the default.policy.users property.

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

The rolling restart action does not work in Kafka Connect when the ssl.client.auth option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set ssl.client.auth to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Known Issues in Spark

Learn about the known issues in Spark the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

The following known issues were identified in this release:

No new known issues in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the RuntimeWarning: Failed to add file message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

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

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

The following services support DDL statements:

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

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

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

CDPD-67517: Spark3 tests are failing if /tmp is mounted as NOEXEC.

Map the tmpdir to a writable path in spark3-conf/spark-defaults.conf using the following steps:

1. Go to the Cloudera Manager.
2. Go to the Spark service.
3. Click the Configuration tab.
4. Select Scope Gateway .
5. Select Category Advanced .
6. Locate the Spark Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-defaults.conf_client_config_safety_valve property.
7. Map the tmpdir to a writable path:

```
spark.driver.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
spark.executor.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
```

8. Enter a Reason for change, and then click Save Changes to commit the changes.
9. Deploy the client configuration.

CDPD-23817: In the upgraded Cluster, the permission of /tmp/spark is restricted due to the HDP configuration hive.exec.scratchdir=/tmp/spark.

If you are using the /tmp/spark directory in the CDP cluster, you must provide the required additional Policies/ACL permissions.

CDPD-22670 and CDPD-23103: There are two configurations in Spark, "Atlas dependency" and "spark_lineage_enabled", which are conflicted. The issue is when Atlas dependency is turned off but spark_lineage_enabled is turned on.

Run Spark application, Spark will log some error message and cannot continue. That can be restored by correcting the configurations and restarting Spark component with distributing client configurations.

CDPD-23007: Mismatch in the Spark Default DB Location. In HDP 3.1.5, hive_db entities have one attribute - 'location' which is configured to the '/managed' path. In fresh install of CDP 7.1.7, hive_db entities now have 2 attributes 'location' configured to '/external' path and 'managedLocation' configured to '/managed' path. In the AM2CM migration (HDP 3.1.5 -> CDP 7.1.7), the 'location' attribute from hive_db entities in HDP 3.1.5 comes unaltered to CDP 7.1.7 and hence maps to '/managed' path.

This issue arises only if you are upgrading from HDP 3.1.5 to CDP 7.1.7. If you are performing a fresh install of CDP 7.1.7, you can ignore this issue.

None

CDPD-217: The Apache Spark connector is not supported

The old *Apache Spark - Apache HBase Connector* (shc) is not supported in CDP releases.

Use the new HBase-Spark connector shipped in CDP release.

CDPD-3038: Launching pyspark displays several HiveConf warning messages

When pyspark starts, several Hive configuration warning messages are displayed, similar to the following:

```
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of name hive.meta
store.runworker.in does not exist
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of
name hive.masking.algo does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of
name hive.metastore.runworker.in does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of na
me hive.masking.algo does not exist
```

These errors can be safely ignored.

Known Issues in Sqoop

Learn about the known issues in Sqoop, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version

If you have upgraded to CDP PvC Base 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to CDP PvC Base 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see [Troubleshooting Apache Sqoop issues](#).

CDPD-44431: Using direct mode causes problems

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the `--direct` option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the `sqoop.enable.deprecated.direct` property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through `-Dsqoop.enable.deprecated.direct=true`.

CDPD-3089: Avro, S3, and HCat do not work together properly

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

Parquet columns inadvertently renamed

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.

PARQUET-99: Importing Parquet files might cause out-of-memory (OOM) errors

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.

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 in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager

Cloudera Manager does not display a Log Files menu for SMM UI role (and SMM UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by SMM UI is not supported by Cloudera Manager.

View the SMM UI logs on the host.

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in SMM safety valve in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Add the following value for bootstrap servers>Save Changes > Restart SMM.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

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

KNOX does not send neither the Authentication header nor the `hadoop.auth` cookie. Because of this, the SMM UI returns an HTTP 401 response and sets the `www-authenticate: Negotiate` header. After this, KNOX still does not send the Authentication header. This results in the SMM UI being inaccessible through Knox.

You can access the SMM UI directly in Cloudera Manager at `Clusters SMM Streams Messaging Manager Web UI` and log in using the Kerberos username and password.

CDPD-82560: Streams Messaging Manager known unprotected endpoint

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

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

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.snapshot` 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 in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

CDPD-11079: Blacklisted topics appear in the list of replicated topics

If a topic was originally replicated but was later excluded for replication, 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, this topic will be included in the response. Additionally, the excluded topic will also be visible in the SMM UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

CDPD-30275: SRM may automatically re-create deleted topics on target clusters

If `auto.create.topics.enable` is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with `srm-control`. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1]
```

2. Wait until SRM is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

Limitations

SRM cannot replicate Ranger authorization policies to or from Kafka clusters

Due to a limitation in the Kafka-Ranger plugin, SRM cannot replicate Ranger policies to or from clusters that are configured to use Ranger for authorization. If you are using SRM to replicate data to or from a cluster that uses Ranger, disable authorization policy synchronization in SRM. This can be achieved by clearing the Sync Topic Acls Enabled (`sync.topic.acls.enabled`) checkbox.

SRM cannot ensure the exactly-once semantics of transactional source topics

SRM data replication uses at-least-once guarantees, and as a result cannot ensure the exactly-once semantics (EOS) of transactional topics in the backup/target cluster.



Note: Even though EOS is not guaranteed, you can still replicate the data of a transactional source, but you must set `isolation.level` to `read_committed` for SRM's internal consumers. This can be done by adding `[***SOURCE_CLUSTER_ALIAS***]->[***TARGET_CLUSTER_ALIAS***].consumer.isolation.level=read_committed` to the Streams Replication Manager's Replication Configs SRM service property in Cloudera Manager. The `isolation.level` property can be set on a global connector or replication level. For example:

```
#Global connector level
connectors.consumer.isolation.level
=read_committed
#Replication level
uswest->useast.consumer.isolation.level=read_committed
```

SRM checkpointing is not supported for transactional source topics

SRM does not correctly translate checkpoints (committed consumer group offsets) for transactional topics. Checkpointing assumes that the offset mapping function is always increasing, but with transactional source topics this is violated. Transactional topics have control messages in them, which take up an offset in the log, but they are never returned on the consumer API. This causes the mappings to decrease, causing issues in the checkpointing feature. As a result of this limitation, failover operations for transactional topics is not possible.

Known Issues in YARN 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 in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12**COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."**

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

None.

COMPX-4644: Queue capacity rounding problem when configuration is initially set via YARN

When setting the capacity scheduler configuration through the YARN/Cloudera Manager configuration, there may be capacity values that use multiple decimal places. This results in rounding/floating point precision discrepancies in the UI when trying to validate that all sibling capacities equal 100%. The UI looks like all the numbers add up to 100, but the validation still displays an error and does not allow to save the capacities. It is also observed that the capacity is being calculated as, for example, 99.9999999991 in the backend.

- Create queues within the UI, or
- Ensure that capacities configured through the Capacity Scheduler safety valve do not have more than one decimal place.

20202 Database migration after enabling opt-in migration

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to CDP 7.1.9, forcing migration from H2 to PostgreSQL database.
- Upgrade to CDP 7.1.9 CHF2, moving back to H2 database.
- Upgrade to CDP 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version CDP 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

CDPD-56559: MapReduce jobs can intermittently fail during a rolling upgrade.

During a rolling upgrade between CDP versions 7.1.8 and 7.1.9, MapReduce jobs may fail with message, `RuntimeException: native snappy library not available`. Although the native Snappy compression library is not loaded, a checkmark displays to indicate that the Snappy compression library is loading for NodeManagers that are pending upgrades. This causes the MapReduce jobs that are associated with the NodeManagers to fail. After the upgrade, the jobs work as expected. This issue only impacts rolling upgrades from before CDP 7.1.9 to a higher version.

None.

COMPX-12021 Queue Manager configurations on Scheduler Configuration page are not working

When setting the following properties on the YARN Queue Manager UI, the properties are set in the `capacity-scheduler.xml` file which does not have any effect on YARN. The properties need to be set in the `yarn-site.xml` file, which does not happen when you set them through YARN Queue Manager.

- "Maximum Application Priority" – `yarn.cluster.max-application-priority`
- "Enable Monitoring Policies" – `yarn.resourcemanager.scheduler.monitor.enable`
- "Monitoring Policies" – `yarn.resourcemanager.scheduler.monitor.policies`



Note: You can also set this property on the YARN configuration page in Cloudera Manager as "Capacity Scheduler Preemption".

- "Preemption: Observe Only" – `yarn.resourcemanager.monitor.capacity.preemption.observe_only`

- "Preemption: Monitoring Interval (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.monitoring_interval"`
 - "Preemption: Maximum Wait Before Kill (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.max_wait_before_kill"`
 - "Preemption: Total Resources Per Round" – `"yarn.resourcemanager.monitor.capacity.preemption.total_preemption_per_round"`
 - "Preemption: Over Capacity Tolerance" – `"yarn.resourcemanager.monitor.capacity.preemption.max_ignored_over_capacity"`
 - "Preemption: Maximum Termination Factor" – `"yarn.resourcemanager.monitor.capacity.preemption.natural_termination_factor"`
 - "Enable Intra Queue Preemption" – `"yarn.resourcemanager.monitor.capacity.preemption.intra-queue-preemption.enabled"`
1. In Cloudera Manager, select the YARN service.
 2. Click the Configuration tab.
 3. Search for `yarn-site.xml`.
 4. Under YARN Service Advanced Configuration Snippet (Safety Valve) for `yarn-site.xml`, add the corresponding parameter and value you need.
 5. Click Save Changes.
 6. Restart the YARN services.

COMPX-13177: QueueManager webapp requests fail with 'HTTP ERROR 400

java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA'

Products:

- Cloudera Manager for CDP Private
- Cloud Base Cloudera Manager for CDP Public Cloud

Context:

- Centos 7.8 and Redhat 7.8 operating systems, when FIPS support is enabled.

Problem:

- When attempting to display the Yarn Queue Manager interface, Cloudera Manager displays an error: "HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA".

1. Edit the file `[/etc/default/cloudera-scm-server]`
2. Around line 28, modify the line that starts with

```
#export CMF_OVERRIDE_TLS_CIPHERS= . . . .
```

3. Remove the comment mark #.
4. Remove all ciphers with "3DES" in the name.
5. Save the file.
6. Restart the Cloudera Manager Server service.

Third-party applications do not launch if MapReduce framework path is not included in the client configuration

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the `mapreduce.application.framework.path` property is set to the appropriate value, but third-party applications with their own configurations does not launch.

Set the `mapreduce.application.framework.path` property to the appropriate configuration for third-party applications.

JobHistory URL mismatch after server relocation

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated the mapred-site.xml file that references the correct JobHistory Server.

YARN cannot start if Kerberos principal name is changed

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN does not start. In such cases, the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the znodel and restart the YARN service.
- Use the reset ZK ACLs command. This also sets the znodels below /rmstore/ZKRMStateRoot to world:anyone:cdwra which is less secure.

Queue Manager does not open on using a custom user with a default Kerberos principal

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager process_username property matches the YARN process_username property.

Known Issues in Zeppelin

Learn about the known issues in Zeppelin the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

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.

BUG-125263: Zeppelin service move fails on clusters upgraded from HDP3.1.5

Resolve the circular symlink issue on the Zeppelin node by linking the conf directory to a new directory under /etc/zeppelin:

- # mkdir -p /etc/zeppelin/<version>/0
- # rm /usr/hdp/<version>/zeppelin/conf
- # ln -s /etc/zeppelin/<version>/0 /usr/hdp/<version>/zeppelin/conf

Where version is the HDP version. For example, 7.1.x-yyy. Restart the Zeppelin server in Ambari.

CDPD-3090: Due to a configuration typo, functionality involving notebook repositories does not work

Due to a missing closing brace, access to the notebook repositories API is blocked by default.

From the CDP Management Console, go to Cloudera Manager for the cluster running Zeppelin. On the Zeppelin configuration page (Zeppelin serviceConfiguration), enter shiro urls in the Search field, and then add the missing closing brace to the notebook-repositories URL, as follows:

```
/api/notebook-repositories/** = authc, roles[{{zeppelin_admin_group}}]
```

Click Save Changes, and restart the Zeppelin service.

CDPD-2406: Logout button does not work

Clicking the Logout button in the Zeppelin UI logs you out, but then immediately logs you back in using SSO.

Close the browser.

OPSAPS-59802: Zeppelin and Livy roles should be co-located on the same host.

When installing or upgrading to CDP Private Cloud Base, you must co-locate all Zeppelin and Livy roles on the same cluster host due to an issue with certificate generation.

Known Issues in ZooKeeper

Learn about the known issues in ZooKeeper the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 12.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 12

There are no new known issues identified in the previous releases.

Cumulative hotfix 11

Know more about the cumulative hotfix 11 for CDP 7.1.9 SP1. This cumulative hotfix was released on September 10, 2025.

Cloudera Runtime 7.1.9.1059-4 (Cumulative Hotfix 11) download URL:**Parcel Repository Location**

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1059/parcels/
```

Replication Manager

Before you use Replication Manager with CDP Private Cloud Base 7.1.9 SP1 CHF11 and Cloudera Manager 7.13.1.400, you must be aware that certain Cloudera Replication Manager fixes are not included in this combination. While general usage is supported, specific features or use cases might be impacted if they rely on the missing fixes.

To ensure full feature compatibility and to access the latest fixes, Cloudera recommends upgrading to Cloudera Manager 7.13.1.500 or higher versions when using CDP Private Cloud Base 7.1.9 SP1 CHF11.

What's new in Cloudera Runtime 7.1.9 SP1 CHF 11

Understand the functionalities and improvements to features of components in Cloudera Runtime 7.1.9 SP1 CHF 11.

What's new in Hue**Enhanced file extension controls for Hue file upload**

Previously, Hue allowed only restricting specific file extensions across all configured filesystems to enhance security.

Hue now allows you to control both allowed and restricted file type extensions for Hue file uploads. This improves flexibility and enhances the overall user experience.

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

Python 3.11 support in Impala-shell

The Impala shell now supports Python 3.11.

Fixed issues in 7.1.9 SP1 CHF 11

Know more about the cumulative hotfix 11 for 7.1.9 SP1.



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

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p1059.70393529.

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

Previously, in the Hue left assist panel, only the initial 5,000 tables were loaded because the default value of `max_catalog_sql_entries` was set to 5,000. Increasing this configuration in Hue Safety Valve only allowed loading up to 10,000 tables due to the `max_rows` parameter being hardcoded to 10,000.

This issue is now resolved by fixing the hardcoded limitation. You can now load more than 10,000 tables in the left assist panel.

CDPD-84277: Query Processor cleanup failures on MySQL due to DELETE syntax errors

Previously, the cleanup DELETE operations in Query Processor were failing on MySQL versions 8.0.25 to 10.3.32 due to the unsupported use of the **USING** clause in **DELETE** statements. The legacy SQL syntax DELETE or USING caused syntax errors and prevented proper execution of cleanup tasks. This issue is now resolved by replacing the **USING** clause with proper **JOIN** syntax in the **DELETE** queries.

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

Previously, in the Hue 7.1.9 SP1 CHF4 release, 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 and Table browser without affecting default system performance. 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.

CDPD-55789: Memory leak in NSSummary in Recon

Previously, the memory in the NSSummary map used to leak when multiple files and directories were deleted and some orphan directory links were unlinked from the NSSummary tree, but, the orphan links could remain in the map forever. This fix cleans up the entries including the deleted files and directories and also the orphan directory links from both the NSSummary map and the NSSummary tree after the directories and files are completely deleted from Ozone, including the `deletedDirTable` table.

Apache Jira: [HDDS-8565](#)

CDPD-87864: Retrigger of Recon NSSummary tree build is not controlled with central flag

Previously, the retriggering of the Recon NSSummary tree was called from multiple places in the source code and was not controlled using a common flag. This issue is now fixed by controlling the retriggering using a central flag.

Apache Jira: [HDDS-13443](#)

CDPD-89346/CDPD-89159: Enhanced join strategy selection for large clusters

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

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

Apache Jira: [IMPALA-14263](#)

Stable catalog cache after INVALIDATE METADATA

Previously, the `INVALIDATE METADATA <table>` command on tables created outside Impala caused tables to be temporarily dropped from the catalog cache.

This issue is now resolved by ensuring that the `INVALIDATE METADATA` command fetches the latest valid event ID from the Hive Metastore and assigns it to the table.

Apache Jira: [IMPALA-12712](#)

CDPD-88652: Catalogd HA warm failover

In a Catalog high availability (HA) setup, a failover cause service delays because the standby Catalogd has a cold cache so need to load the metadata again which takes time. This is particularly noticeable when metadata of lots of the tables have been loaded in the active Catalogd.

This issue is now resolved by allowing the standby Catalogd to keep its metadata cache ready, including a new method, added in [IMPALA-14074](#), to pre-load metadata of critical tables and keep the tables up-to-date. As a result, when a failover happens, the new active Catalogd can serve requests on critical tables right away without any delays.

The fix includes the following JIRA changes:

1. • [IMPALA-14074](#) – Adds the following flags to enable metadata cache warmup in Catalogd for critical tables, ensuring that the standby is ready for failover:

- `--warmup_tables_config_file` – Specifies a configuration file with the tables whose metadata is to be preloaded. The file can be on local storage or remote storage, such as HDFS. Each line can list a fully qualified table or a wildcard for all tables under a database, for example, `tpch.*`. Tables are loaded in the order listed in the file.

Example:

```
--warmup_tables_config_file=file:///opt/impala/
warmup_table_list.txt
--warmup_tables_config_file=hdfs:///tmp/
warmup_table_list.txt
```

- `--keeps_warmup_tables_loaded` – Controls whether the listed tables are automatically reloaded after invalidation. By default, this is set to false.



Important: Avoid setting this flag to true if the Catalogd heap size is not large enough to cache all the warmup tables. Otherwise, tables might be repeatedly loaded and invalidated, causing memory pressure.

- `--catalogd_ha_reset_metadata_on_failover` – Must be set to false to ensure that the standby Catalogd maintains a warm metadata cache for a faster switch during failover.
2. [IMPALA-14227](#) – Adds a waiting time during HA failover to ensure the new active Catalogd applies all the pending HMS events before serving requests.
 3. [IMPALA-12876](#) – Exposes the Catalog version and loaded timestamps in query profiles to help debug stale metadata issues.

Apache Jira: [IMPALA-14074](#), [IMPALA-14227](#), [IMPALA-12876](#)

New configuration to disable block location retrieval

Impala is designed to optimize reads by retrieving file block locations from storage systems like HDFS and Ozone, and scheduling read operations on the same nodes where the data resides. In on-premise environments, Impala nodes are typically colocated with storage nodes. However, in scenarios where they are not colocated, retrieving the block location information is a resource-intensive operation.

This issue is now fixed by introducing new Hadoop configuration properties that allow you to disable the retrieval of block location information during table loading. Disabling the retrieval of block location information can be configured in the following ways:

- Globally, using the Hadoop Safety Valve for core-site.xml advanced configuration snippet.

```
'impala.preload-block-locations-for-scheduling': 'false'
```

- By using a Filesystem scheme to disable it for a specific type of storage.

```
'impala.preload-block-locations-for-scheduling.scheme.hdfs':  
'false'
```

- By authority, to disable block location retrieval for a specific endpoint when multiple storage systems are configured with the same scheme.

```
'impala.preload-block-locations-for-  
scheduling.authority.mycluster': 'false'
```

This disables block location retrieval only for URIs as `hdfs://mycluster/warehouse/tablespace/...`

Apache Jira: [IMPALA-14138](#)

Fix performance regression on Catalog operations when Catalogd HA is enabled

Previously, when Catalogd HA was enabled, performance regression occurred because a lock contention issue blocked new requests. This was caused by an additional check in each request that unnecessarily acquired an exclusive lock, impacting concurrent performance.

This issue is now resolved by optimizing the check to not depend on the exclusive lock. This change improves the concurrent performance of Catalogd in HA environments.

Apache Jira: [IMPALA-14220](#)

CDPD-87722: Merge/Update jobs fail with vertex errors after upgrading to 7.1.9

In CDP 7.1.9, improvements for update and delete operations were introduced to write data directly to table locations instead of using a staging directory. However, after upgrading to CDP 7.1.9, merge, update, and delete queries might occasionally fail with the following error:

```
java.lang.RuntimeException: Hive Runtime Error while closing  
operators: Index 17 out of bounds for length 17  
    at  
    org.apache.hadoop.hive ql.exec.tez.TezProcessor.initializeAndRunProcessor(TezPr  
    at  
    org.apache.hadoop.hive ql.exec.tez.TezProcessor.run(TezProcessor.java:252)  
    at  
    org.apache.tez.runtime.LogicalIOProcessorRuntimeTask.run(LogicalIOProcessorRunt  
    .....  
    Caused by: java.lang.RuntimeException: Hive Runtime  
Error while closing operators: Index 17 out of bounds for length  
17  
    at  
    org.apache.hadoop.hive ql.exec.tez.ReduceRecordProcessor.close(ReduceRecordPro  
    .....
```

```

    Caused by: java.lang.ArrayIndexOutOfBoundsException:
    Index 17 out of bounds for length 17
    at org.apache.hadoop.hive.ql.exec.FileSinkOperator
    $FSPaths.closeWriters(FileSinkOperator.java:258)
    at
    org.apache.hadoop.hive.ql.exec.FileSinkOperator.closeOp(FileSinkOperator.java:1
    at
    org.apache.hadoop.hive.ql.exec.Operator.close(Operator.java:731)
    at
    org.apache.hadoop.hive.ql.exec.Operator.close(Operator.java:755)
    at
    org.apache.hadoop.hive.ql.exec.tez.ReduceRecordProcessor.close(ReduceRecordProc

```

This issue is now fixed.

CDPD-88850/CDPD-89412: Compaction cleaner prematurely cleans up deltas

Previously, queries that ran for a long time failed after a compaction because the compaction cleaner prematurely removed delta files. This was a regression introduced by a change in HIVE-23107, which removed a mechanism used to prevent the deletion of files still needed by running queries.

This issue is now resolved by ensuring the compaction cleaner waits for all previous transactions to commit before it begins cleaning delta files.

Apache Jira: [HIVE-24291](#), [HIVE-23107](#)

CDPD-88851: Backward compatibility for Hive Metastore schema

An update from HIVE-23107 caused a problem when you used different versions of Hive Metastore services in the same environment. This was because the newer version removed a database table called `min_history_level` that was required for older versions to function correctly.

This issue is now resolved by a new change that keeps the `min_history_level` table.

Apache Jira: [HIVE-24403](#)

CDPD-88487: Queries with subqueries failing compilation

After upgrading to CDP 7.1.9 SP1 CHF10, some queries that included subqueries failed during compilation. The queries failed with the `UnsupportedOperationException` error.

This issue is now resolved by correcting the compilation failure that occurred in the `SharedWorkOp` timer.

CDPD-89289: Backport RATIS-1884 to 7.1.9 SP1

Previously, a regression occurred in 7.1.9 SP1 CHF 9 and 7.1.9 SP1 CHF 10 where the Ozone Manager log listed the `RUNNING` retry cache entry of leader should be pending Ratis error message.

This issue is now resolved.

Apache Jira: [RATIS-1884](#)

CDPD-23034: Port COMPLEX-17261 to 7.1.9.1000

The `yarn.nodemanager.disk-health-checker.min-free-space-per-disk-watermark-high-mb` new optional configuration value is now created. This value refers to the minimum space in megabytes that must be available on a bad disk for it to be marked as good. This value must not be less than the `yarn.nodemanager.disk-health-checker.min-free-space-per-disk-mb` value. If it is less than the `yarn.nodemanager.disk-health-checker.min-free-space-per-disk-mb` value or not set, the same value is used as the `yarn.nodemanager.disk-health-checker.min-free-space-per-disk-mb` value. This setting applies to both the `yarn.nodemanager.local-dirs` and `yarn.nodemanager.log-dirs` directories.

CDPD-22902: 7.1.9.1000: RM UI2 cannot show logs for non-MapReduce jobs with MR ACLs enabled

Previously, container log ACL checks failed for non-MapReduce jobs.

This issue is now resolved by applying the MapReduce ACLs only if the container belongs to a MapReduce job

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

CVE-2017-7525	jackson-mapper-asl
CVE-2019-10172	jackson-mapper-asl
CVE-2025-52434	Apache Tomcat
CVE-2025-52520	Apache Tomcat
CVE-2025-53506	Apache Tomcat
CVE-2024-13009	Eclipse Jetty
CVE-2024-8184	Eclipse Jetty

Behavioral changes in Cloudera Runtime 7.1.9 SP1 CHF 11

Know more about the behavioral changes in cumulative hotfix 11 for 7.1.9 SP1.

HBase

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](#)

Impala

Summary:

Batched Insert events for Impala

Previous Behavior: Impala initiated insert statements using a single API call for each partition. This could become a performance bottleneck when inserting data into tables with a large number of partitions.

New Behavior: This change speeds up computing file checksums during INSERT operations, especially for tables with many files and partitions. This is achieved by leveraging a new batch insert API in the Hive Metastore to process events more efficiently.

Apache Jira: [IMPALA-14051](#)

Summary:

Support for parquet PageHeaderV2

Previous Behavior: The parquet scanner did not support reading or processing of V2 Parquet DataPageHeader.

New Behavior: The parquet scanner now supports reading of V2 Parquet DataPageHeader.

Apache Jira: [IMPALA-6433](#)

Summary:

Faster checksum computation during INSERT

Previous Behavior: During INSERT operations, most of the time was spent fetching file checksums, which was a slow process. The system acquired a table lock before performing these checksum calculations, causing other operations that required the lock to be blocked.

New Behavior: Impala now collects file checksums and ACID directory paths before acquiring the table lock. This change ensures that the time-consuming process of fetching checksums does not block other operations. As a result, computing file checksums during INSERT operations with many files is now faster.

Apache Jira: [IMPALA-12162](#)

Known issues in 7.1.9 SP1 CHF 11

The following section lists the known issues and limitations, the areas of impact, and workaround for each Runtime component.

Known Issues in Atlas

Learn about the known issues in Atlas the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

CDPD-93307: Hive import logs are truncated during bulk import

When running the import-hive.sh script to perform bulk Hive metadata imports, the logging process fails to retain the expected history. The import-hive.log file rotates prematurely when it reaches 10 MB and retains only one backup file (import-hive.log.1), instead of the expected 100 MB file size and 20 backup files.

Supply a custom log4j configuration file with the correct structure by using the ATLAS_OPTS environment variable before running the import script:

```
<?xml version="1.0" encoding="UTF-8" ?>
<!DOCTYPE log4j:configuration SYSTEM "log4j.dtd">
<log4j:configuration xmlns:log4j="http://jakarta.apache.org/lo
g4j/">
  <appender name="console" class="org.apache.log4j.ConsoleAppen
der">
    <param name="Target" value="System.out"/>
    <layout class="org.apache.log4j.PatternLayout">
      <param name="ConversionPattern" value="%d %-5p - [%t:
%x] ~ %m (%C{1}:%L)%n"/>
    </layout>
  </appender>
  <appender name="FILE" class="org.apache.log4j.RollingFileApp
ender">
    <param name="File" value="${atlas.log.dir}/${atlas.log.f
ile}"/>
    <param name="Append" value="true"/>
    <param name="maxFileSize" value="100MB"/>
    <param name="maxBackupIndex" value="20"/>
    <layout class="org.apache.log4j.PatternLayout">
      <param name="ConversionPattern" value="%d %-5p - [%t:
%x] ~ %m (%C{1}:%L)%n"/>
    </layout>
  </appender>
  <logger name="org.apache.atlas" additivity="false">
    <level value="info"/>
    <appender-ref ref="FILE"/>
  </logger>
  <logger name="org.apache.kafka.common.config.AbstractConfig"
additivity="false">
    <level value="error"/>
    <appender-ref ref="FILE"/>
  </logger>
```

```
<root>
  <priority value="warn"/>
  <appender-ref ref="FILE"/>
</root>
</log4j:configuration>
```

Export the ATLAS_OPTS variable to point to your custom file before running the import script:

```
export ATLAS_OPTS="-Dlog4j.configuration=file:/path/to/custom-hi
ve-import-log4j.xml"
./import-hive.sh
```

CDPD-87670: [Atlas UI] Apache Atlas Glossary becomes unresponsive when the page size is set to 50

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

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-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-6565: Issue with ddlQueries in Atlas for Hive/Impala tables created in default/non-default db using DAS/HUE respectively

ddlQueries are not created for Impala origin table. ddlQueries created for Impala/Hive CTAS tables have difference in names as compared to same operation for Hive table. The Impala table has the query in the ddlQueries entity name.

None

CDPD-55301: ddlQueries and ALTERNATIVE_* lineage missing for Spark tables created through spark3-shell

The ddlQueries and ALTERNATIVE_* lineage missing for Spark tables created using spark3-shell.

None

CDPD-56085: 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

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*.

None

CDPD-58554: Discard audits of specific classification, label, and business metadata

Support to control audits for specific classification, label, and business metadata is not present in the Custom Audit Filters feature.

None

CDPD-40346: The ddlQueries and ALTERNATIVE_ADDCOLS lineage missing for Impala tables.

The ALTERNATIVE_ADDCOLS lineage has some issue when an Impala table is altered and the corresponding lineage is not created.

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

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

None

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

Adding Japanese labels or user defined properties results in the error message: “Invalid label: ###, label should contain alphanumeric characters, _ or -”

None

CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties

Entities, such as hdfs_path, are not returned by the search when their indexable properties are searched using partial search terms made of Japanese or Chinese characters. Only exact matches return results when searching the indexed properties made of Japanese or Chinese characters.

None

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

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

None.

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

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

None.

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

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

None.

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

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

None.

CDPD-43772: Performance issues with Atlas service

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

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

OPSAPS-67783: During rolling upgrade one among two Atlas server failed to start but Cloudera Manager considered as success

Cloudera Manager marks the Execute command Start on service Atlas-1 as a success even when Atlas service had failed to start successfully. In such cases, Atlas logs give the exact reason for Atlas start-up failure.

None

CDPD-19358: "IsIndexable" and "isOptional" value of a typedef's attribute is modified post migration.

In HDP-265, falcon_feed_creation has "stored-in" attribute which has IsIndexable value set to True.

Post migration , "stored-in" is moved to relationshipAttributeDefs and has IsIndexable value set to False.

Similarly, in HDP-265 the table attribute of `hive_storage_desc` isOptional is set to true. Post migration, isOptional set to False.

None.

CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command

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

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

In converting content from Navigator to Atlas, the create time for Hive tables is not moved to Atlas.

None

CDPD-10574: Suggestion order doesn't match search weights

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

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) does not 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 the metadata collection from HMS. For example, if an ALTER TABLE operation in Spark that is changing a table name and is reporting 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 "@" does not 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, and use 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 does 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 logout 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 CDP.

None.

CDPD-1664: Guest users are redirected incorrectly

Authenticated users logging in to Atlas are redirected to the CDP 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: The 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, to add metadata tags to relationships that you wanted to make sure were unique in the system, you can design the relationship attribute with the property IsUnique equal true. However, in this release, the IsUnique attribute is not enforced.

None.

DOCS-13759: Tag Propagation stops after a certain depth while the lineage is being extended

When a tag is added to an entity at timestamp T1, the entities along the lineage to which the tag must be propagated is calculated at T1. Before tag propagation completes, if the lineage is extended, tag does not propagate to the entities in the extended lineage.

CDPD-41142: When a Kafka console consumer group is run, more than one update audits are seen

After running the console consumer with a consumer group, verify the consumer group entity created, along with the metrics and notifications for the consumer group and topic. The expected result can be: one ENTITY_CREATE audit and one ENTITY_UPDATE audit. But more than one ENTITY_UPDATE audits are seen.

CDPD-40165: Two audits are created for SPARK CTAS table

When following Spark queries are fired:

```
spark.sql("create table table1(id int)")
spark.sql("create table table2 as select * from table1")
```

HMS sends "ENTITY_CREATE" and "ENTITY_FULL_UPDATE_V2".

The extra ENTITY_FULL_UPDATE_V2 message received from HMS is sent as a part of ALTER_TABLE_ADDCOLS event from the HMS Hook side. This behaviour is observed only when the queries are run from Spark SQL and not when run the same queries are run from Beeline.

CDPD-39197: Debug metrics returns empty data

When debug metrics is enabled and some operations are performed, the response is empty

CDPD-36495: Updating legacyAttribute from False to True resets the initially created relationshipAttributes values

Creating types, entities, and to start, you must set the relationship with is_legacy_attribute value as **False**.

Later, update the value relationshipDef is_legacy_attribute to **True**.

For the entities that were created before updating the is_legacy_attribute to **True**, the relationship Attributes value is reset.

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-29409: Hive import: Suggestion suggests entity which is deleted.

Suggestions suggests tables of a database, which is a deleted entity.

CDPD-25152: Tag propagation through deferred actions consumes additional time as compared to default flow

The additional time might be due to the small overhead added to create / update task vertex and which is run in the background. This also depends on number of tasks queued to be executed.

CDPD-42954: Zeppelin notebook fails after enabling Atlas-HDFS hook

The Zeppelin notebooks are failing with errors after enabling Atlas-HDFS hook in the CDP cluster.

When the below properties are set for atlas-client.properties in Cloudera Manager:

- atlas.jaas.KafkaClient.option.keyTab
- atlas.jaas.KafkaClient.option.principal

Along with adding the properties in /etc/atlas/conf/atlas-application.properties, Cloudera Manager also adds these properties to atlas-application.properties for other services (like Spark).

Adding these properties interferes with the normal flow of the services (like Spark)

To enable HDFS lineage feature, instead of setting these properties through Cloudera Manager, users can manually add the properties directly in /etc/atlas/conf/atlas-application.properties

CDPD-10576: Deleted Business Metadata attributes appear in Search Suggestions

Atlas search suggestions continue to show Business Metadata attributes even if the attributes have been deleted.

None.

CDPD-39427: [HDFS Lineage]: When the input is a directory in case of put/copyfromLocal/cp/mv commands, lineage is not created even though the script succeeds.

When Source is a directory and target is a directory which is already present in Atlas, the command succeeds and inserts the data in the desired location, but lineage is not created.

DOCS-13760: System Attributes search, __classificationNames: Search with parent tag does not return entities associated to its children tags

System attribute search with __classificationNames = parent_tag returns entities associated to parent_tag only and not entities associated to its children tag.

Workaround: Instead of using system attribute, employ the basic search attribute "classification" which lists entities associated with inherited classifications.

CDPD-35818: Basic search with tag filter provides approximateCount as -1 when there is no match and is 0 otherwise

When the following search operations are performed:

- Faceted search with both tag filter and entity filter

The observed approximateCount is -1.

Here both entity filter and tag filter are present and when there is no match the response received is -1.

- Faceted search with only entity filter

Performing a basic search provides an approximate value of 0 when there is no match.

- Faceted search with only tag filter

Whenever there is tag filter in the query and there is no entity match, the approximateCount is -1 and if the tag filter is not available, the response approximateCount is 0

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export

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

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](#)

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

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

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

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

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

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

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

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

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-88230: Bulk import operations currently do not construct lineage as expected

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

Bulk import operations do not create audits for the imported entities.

Known Issues in Cruise Control

Learn about the known issues in Cruise Control, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11**Rebalancing with Cruise Control does not work due to the metric reporter failing to report the CPU usage metric**

On the Kafka broker, the Cruise control metric reporter plugin may fail to report the CPU usage metric.

If the CPU usage metric is not reported, the numValidWindows in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.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 the [Moving a Host Between Clusters](#).



Note: Cluster nodes affected by this issue are not displayed as unhealthy.

Known Issues in Hue

Learn about the known issues in Hue, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11**OPSAPS-75134: LDAP and Kerberos dual authentication fails with HiveOnTez in HTTP Transport Mode**

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 go to 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
```

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

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

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark queries in a batch mode launches a shell script Oozie job from Hue and this is not supported on a secure cluster.

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

OPSAPS-69659: Hue service fails on restart with "Unable to find pycogp2 2.5.4" error

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.

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

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

None.

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

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

None.

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

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

None.

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

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

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

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

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

CDPD-43293: Unable to import Impala table using Importer

Creating Impala tables using the Hue Importer may fail.

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

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

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

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

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

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

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

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

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

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

Hue 4 web interface uses the unsafe-inline directive in its CSP header. As a result, the application server does not set the CSP header in its HTTP responses, and therefore does not benefit from the additional protection against potential cross-site scripting issues and other modern application vulnerabilities which a properly configured CSP may provide. This could lead to application vulnerability.

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

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

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

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

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

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

For example:

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

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

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

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

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

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

Invalid S3 URI error while accessing S3 bucket

The Hue Load Balancer merges the double slashes (//) in the S3 URI into a single slash (/) so that the URI prefix `"/filebrowser/view=S3A:/"` is changed to `"/filebrowser/view=S3A:/"`. This results in an error when you try to access the S3 buckets from the Hue File Browser through the port 8889.

The Hue web UI displays the following error: “Unknown error occurred”.

The Hue server logs record the “ValueError: Invalid S3 URI: S3A” error.

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.

You should be able to load the S3 browser from both 8888 and 8889 ports.

Alternatively, you can use the Hue server port 8888 instead of the load balancer port 8889 to resolve this issue.

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.

- Restart the Hue Load Balancer.

CDPD-17956, CDPD-14167: Setting idle session timeout for Hue does not work when the cluster is secured using Knox SSO

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

None

Query history displays "Result Expired" message

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

None.

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

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

Hue limitation after upgrading from CDH to CDP Private Cloud Base

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

Hue UI is blank upon login after upgrading to CDP 7.1.7 from CDH 6

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

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

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



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

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

Unsupported features

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

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

```
[desktop]
[spark]
```

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

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

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

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

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

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

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

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



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

Known Issues in HDFS

Learn about the known issues in HDFS, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

There are no known issues identified in the release before Cloudera Runtime 7.1.9 SP1 CHF 11.

Known Issues in Hive

Learn about the known issues in Hive, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

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

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

The following services support DDL statements:

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

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

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

CDPD-60770: Beeline Authentication Issue with Special Characters in Passwords

When LDAP is enabled, users cannot authenticate with Beeline if the password contains a special character. For example, the following string fails:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>;user=user@XXX;password='R3G#xpXyoy1M0Jb1'
```

Use the `-p` parameter to execute the Beeline command:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>; -n user@XXX -p 'R3G#xpXyoy1M0Jb1'
```

CDPD-43769: Inconsistent behavior with CHAR comparisons and string literals

In Hive, 'char' data to 'string' data comparison is considered as equal, but 'char' data to 'varchar' data comparison is considered as not equal if the data has trailing spaces. For example,

```
create table test(colchar char(10), colvarchar varchar(20), colstring string);
insert into test values ('a ', 'a ', 'a ');

select count(*) from test where colchar = colstring;
Output returns 1 because char and string data is considered as equal

select count(*) from test where colchar = colvarchar;
Output returns 0 because char and varchar data is considered as not equal
```

It is recommended that you use 'varchar or string' data type instead of 'char' data type for any character or string representation.

CDPD-42726: Wrong results for agg queries from stats after running multi-insert query on managed table

Multiple inserts at the same time into the same partition results in invalid stats. For example:

```
create table source(p int, key int, value string);
insert into source(p, key, value) values (101,42,'string42');

create table stats_part(key int,value string) partitioned by (p int);
from source
insert into stats_part select key, value, p
insert into stats_part select key, value, p;

select count(*) from stats_part;
```

```
-- In this scenario, StatsOptimizer helps serving this query be
cause the result should be rowNum of the partition p=101.
-- The result is 1, however, the result should be 2.
```

Rewrite the multi-insert statement by merging the branches inserting into the same partition to one branch.

CDPD-57574: Query execution fails due to NullPointerException in DagUtils.setupQuickStart

Hive queries may fail at runtime with a NullPointerException while executing the Tez graph at DagUtils.setupQuickStart method. A part of the stack trace from the error is shown below:

```
ERROR : Failed to execute tez graph.
java.lang.NullPointerException: null at org.apache.hadoop.hive.
ql.exec.tez.DagUtils.setupQuickStart(DagUtils.java:1724)
```

The error usually occurs from an invalid plan that has cycles created by semi-join or map-join edges.

Disable some of the optimizations in order to remove the cycles from the plan. It is sufficient to set one of the following properties:

- set hive.tez.dynamic.partition.pruning=false;
- set hive.tez.dynamic.semijoin.reduction=false;
- set hive.auto.convert.join=false;

To avoid any undesired performance regressions, disable the optimization(s) only for specific queries and not globally (through Cloudera Manager) for the whole workload or cluster.

CDPD-45370: Cartesian join fails with an empty table

When the configuration hive.tez.cartesian-product.enabled is set to true, query fails if one of the tables involved in a Cartesian join has zero rows. This results in an IllegalArgumentException with the message Number of items is 0. Should be positive.

To address this issue, set hive.tez.cartesian-product.enabled to false before running the query.

Apache Jira: [\[HIVE-26582\]](#)

CDPD-43107: SemanticException for INSERT INTO statement when Hive Cost-based Optimizer (CBO) is disabled

If you are running the INSERT INTO query and hive.cbo.enable is set to "false", the query fails with a SemanticException. For example,

```
set hive.cbo.enable=false;

CREATE TABLE mytable (
  id INT,
  str STRING
);

INSERT INTO mytable (id, str) VALUES (1, 'a');
```

```
Output:
org.apache.hadoop.hive.ql.parse.SemanticException: 0:0 Expected
2 columns for insclause-0/default@mytable; select produces 1 col
umns. Error encountered near token 'a'
```

This issue occurs only when columns are specified for the target table in the query. INSERT INTO mytable values (1, 'a'); does not result in an exception.

Enable CBO and run the query.

CDPD-43957: HiveServer shuts down during replication due to high resource usage

During Hive replication (both bootstrap and incremental), you may notice that the HiveServer (HS2) shuts down periodically with the following error:

```
java.sql.SQLException: org.apache.hive.jdbc.ZooKeeperHiveClientException: Unable to read HiveServer2 configs from ZooKeeper
    at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.java:265)
    at org.apache.hive.jdbc.HiveDriver.connect(HiveDriver.java:107)
    at java.sql.DriverManager.getConnection(DriverManager.java:664)
    at java.sql.DriverManager.getConnection(DriverManager.java:247)
    at com.cloudera.enterprise.hive3qt.Hive3QueryTool$HiveOperation.execute(Hive3QueryTool.java:682)
    at com.cloudera.enterprise.hive3qt.Hive3QueryTool.main(Hive3QueryTool.java:935)
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Unable to read HiveServer2 configs from ZooKeeper
    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:177)
    at org.apache.hive.jdbc.Utills.configureConnParamsFromZooKeeper(Utills.java:580)
    at org.apache.hive.jdbc.Utills.parseURL(Utills.java:391)
    at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.java:263)
    ... 5 more
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Tried all existing HiveServer2 uris from ZooKeeper.
    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.getServerHosts(ZooKeeperHiveClientHelper.java:132)
    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:172)
    ... 8 more
```

This issue occurs when you replicate at scale with an unbalanced cluster setup that has all the roles running on the Cloudera Manager host. As a result, Cloudera Manager ends up in a bottleneck situation because HS2 crashes, preventing further replication. The logs (/var/log/messages) indicate that there were a large number of Java processes running on the host, which exhausted the host's memory and triggered an Out Of Memory Killer process to stop HS2.

You can restart HS2 and reinitiate replication. However, the replication process may take a longer time to complete.

- Follow these recommendations when setting up your source and target clusters:
 - Ensure that no HS2, Hive metastore (HMS), DataNode, or NameNode roles are running on the Cloudera Manager host.
 - Ensure that you have multiple instances of HS2 and HMS roles on different nodes.
 - It is recommended that you have the HS2 and HMS roles on different nodes than the DataNodes or NameNodes.

These practices increase the possibility of Hive replication at scale completing successfully.

- Limit the scale of Hive replication by temporarily disabling replication policies.

CDPD-40730: Parquet change can cause incompatibility

Parquet files written by the parquet-mr library in CDP 7.1.9, where the schema contains a timestamp with no UTC conversion will not be compatible with older versions of Parquet readers. The effect is that the older versions will still consider these timestamps as they would require UTC conversions and will thus end up with a wrong result. You can encounter this problem only when you write Parquet-based tables using Hive, and tables have the non-default configuration `hive.parquet.write.int64.timestamp=true`.

None.

CDPD-41274: HWC + Oozie issue: Could not open client transport with JDBC Uri

Currently only Spark cluster mode is supported in the Oozie Spark Action with Hive Warehouse Connector (HWC).

Use Spark action in cluster mode.

```
<spark xmlns="uri:oozie:spark-action:1.0">
  ...
  <mode>cluster</mode>
  ...
</spark>
```

CDPD-26556 After an upgrade, querying a CTAS table under certain conditions might throw an exception

If you upgrade your Hive cluster from CDH 6 to CDP 7, create a CTAS table in the CDP cluster from a table you upgraded from CDH, you might see the following exception when you query the new table:

```
class org.apache.hadoop.io.IntWritable cannot be cast to class o
rg.apache.hadoop.hive.serde2.objectinspector.StandardUnionObject
Inspector$StandardUnion
```

This issue involves CDH-based tables having columns of complex types ARRAY, MAP, and STRUCT.

CDPD-23506: OutOfMemoryError in LLAP

Long running spark-shell applications can leave sessions in interactive Hiveserver2 until the Spark application finishes (user exists from spark-shell), causing memory pressure in case of a high number of queries in the same shell (1000+).

You must close spark-shell so that sessions are closed. Add the owner of the database or the tables as a user with read or read/write access to the tables directly.

CDPD-23041: DROP TABLE on a table having an index does not work

If you migrate a Hive table to CDP having an index, DROP TABLE does not drop the table. Hive no longer supports indexes ([HIVE-18448](#)). A foreign key constraint on the indexed table prevents dropping the table. Attempting to drop such a table results in the following error:

```
java.sql.BatchUpdateException: Cannot delete or update a parent
row: a foreign key constraint fails ("hive"."IDX", CONSTRAINT "
IDX_FK1" FOREIGN KEY ("ORIG_TBL_ID") REFERENCES "TBLS ("TBL_ID"
))
```

There are two workarounds:

- Drop the foreign key "IDX_FK1" on the "IDX" table within the metastore. You can also manually drop indexes, but do not cascade any drops because the IDX table includes references to "TBLS".
- Launch an older version of Hive, such as Hive 2.3 that includes IDX in the DDL, and then drop the indexes as described in [Language Manual Indexing](#).

Apache Issue: [Hive-24815](#)

CDPD-17766: Queries fail when using spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive.

Do not specify spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive remotely.

Workaround: specify `principal=hive.server2.authentication.kerberos.principal` as shown in the following syntax:

```
jdbc:hive://<host>:<port>/<dbName>;principal=hive.server2.authentication.kerberos.principal;<otherSessionConfs>?<hiveConfs>#<hiveVars>
```

CDPD-13636: Hive job fails with OutOfMemory exception in the Azure DE cluster

Set the parameter `hive.optimize.sort.dynamic.partition.threshold=0`. Add this parameter in Cloudera Manager (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)

CDPD-10848: HiveServer Web UI displays incorrect data

If you enabled auto-TLS for TLS encryption, the HiveServer2 Web UI does not display the correct data in the following tables: Active Sessions, Open Queries, Last Max n Closed Queries

Known Issues in Impala

Learn about the known issues in Impala, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

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

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

The following services support DDL statements:

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

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

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

CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the `--tls_ciphersuites` startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.

Ensure that the list in the `--tls_ciphersuites` startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

Apache Jira: [IMPALA-14625](#)

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work

If you try to run `impala-shell` on RHEL 9 by setting the default python executable available in `PATH` to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, `IMPALA_PYTHON_EXECUTABLE=python3`.

CDPD-81980: Impala query may fail due to backpressure

Impala queries could fail with a dropped due to backpressure error, even when service queue and memory usage appear normal.

Increase the value of the `--mem_limit` startup option.

CDPD-42958: After upgrading the CDH 7.1.9 from CDH 6.x, under certain conditions you cannot insert data into a table

Under the following conditions, after upgrading from CDH 6.x to CDH 7.1.9 you cannot insert data into a table from Impala:

- On CDH 6.x, you created a database with Impala in a user specified HDFS location.
- Using Hive, you then created a table in the database.

Under these conditions, the database and table are stored in the user-specified HDFS directory. After upgrading, the HDFS directory of the table is read-only for Impala. Consequently, from Impala you cannot insert new data into the table because Impala does not have write permission on the HDFS directory.

Workaround: To resolve this issue, use either one of the following workarounds:

- Using the Ranger Web UI, in the policy repository `cm_hdfs`, grant the user 'impala' write permission on the directory where the table resides.
- Enter the following command to grant write permission to user 'impala' on the HDFS directory where the table resides.

```
hdfs dfs -setfacl -m default:user:impala:rwx <HDFS directory>
```

CDPD-80019: Impala fails to start due to large catalog-update topic

Impala coordinators fail to start if the catalog-update statestore topic exceeds 2GB, hitting the Thrift 0.16.0 message size limit (`thrift_rpc_max_message_size`). This prevents the coordinator from processing the update.

- Enable local catalog mode: `--use_local_catalog=true`
- In Cloudera Manager, go to Impala Configuration search for "Catalog Server Command Line Argument Advanced Configuration Snippet (Safety Valve)" and click 
- Add property name and value:

```
Name : --catalog_topic_mode
Value: minimal
```

- Save changes.

Apache Jira: [IMPALA-12350](#)

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 HDF
S file hdfs://nameservice1/warehouse/tablespace/managed/hive/<da
tabase>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException:
File does not exist: /warehouse/tablespace/managed/hive/<data
base>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

You must use a single Knox node to access Impala UI.

Impala api calls via 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-21828: Multiple permission assignment through grant is not working

None

Problem configuring masking on tables using Ranger

The following Knowledge Base article describes the behavior when we configure masking on tables using Ranger. This configuration works for Hive, but breaks queries in some scenarios for Impala.

For a workaround, see the following Knowledge Base article: [ERROR: "AnalysisException: No matching function with signature: mask\(FLOAT\)" when Impala jobs fail with the following error with signature: mask\(FLOAT\)](#)

IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXXXX"

IMPALA-532: Impala should tolerate bad locale settings

If the LC_* environment variables specify an unsupported locale, Impala does not start.

Add LC_ALL="C" to the environment settings for both the Impala daemon and the Statestore daemon.

Avro Scanner fails to parse some schemas

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

IMPALA-691: Process mem limit does not account for the JVM's memory usage

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the `impalad` daemon.

To monitor overall memory usage, use the `top` command, or add the memory figures in the Impala web UI `/memz` tab to JVM memory usage shown on the `/metrics` tab.

IMPALA-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Remove trailing semicolon from the Avro schema.

IMPALA-1652: Incorrect results with basic predicate on CHAR typed column

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the `RPAD()` function to blank-pad literals compared with `CHAR` columns to the expected length.

IMPALA-1821: Casting scenarios with invalid/inconsistent results

Using a CAST() function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

IMPALA-2005: A failed CTAS does not drop the table if the insert fails

If a CREATE TABLE AS SELECT operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

Drop the new table manually after a failed CREATE TABLE AS SELECT

IMPALA-2422: % escaping does not work correctly when occurs at the end in a LIKE clause

If the final character in the RHS argument of a LIKE operator is an escaped \% character, it does not match a % final character of the LHS argument.

None

IMPALA-2603: Crash: impala::Coordinator::ValidateCollectionSlots

A query could encounter a serious error if includes multiple nested levels of INNER JOIN clauses involving subqueries.

None

IMPALA-3094: Incorrect result due to constant evaluation in query with outer join

An OUTER JOIN query could omit some expected result rows due to a constant such as FALSE in another join clause. For example:

```
explain SELECT 1 FROM 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 |
+---+
- - - - -
```

IMPALA-3509: Breakpad minidumps can be very large when the thread count is high

The size of the breakpoint 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 `-lminidump_size_limit_hint_kb=size` to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then 2 KB per thread after that.) The minidump file can still grow larger than the "hinted" size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

IMPALA-4978: Impala requires FQDN from hostname command on Kerberized clusters

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the `gethostname()` system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a Kerberized cluster.

Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname`, only returns the short name, pass the command-line flag `##hostname=FULLY_QUALIFIED_DOMAIN_NAME` in the startup options of all Impala-related daemons.

IMPALA-6671: Metadata operations block read-only operations on unrelated tables

Metadata operations that change the state of a table, like `COMPUTE STATS` or `ALTER RECOVER PARTITIONS`, may delay metadata propagation of unrelated unloaded tables triggered by statements like `DESCRIBE` or `SELECT` queries.

None

IMPALA-7072: Impala does not support Heimdal Kerberos

None

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.

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.

Known Issues in Iceberg

Learn about the known issues in Iceberg, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

There are no known issues identified in earlier releases.

Known Issues in Kafka

Learn about the known issues in Kafka, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

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

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

The following services support DDL statements:

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

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

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

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker) in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Save Changes > Restart SMM.

The `offsets.topic.replication.factor` property must be less than or equal to the number of live brokers

The `offsets.topic.replication.factor` broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a `GROUP_COORDINATOR_NOT_AVAILABLE` error until the cluster size meets this replication factor requirement.

None

Requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true

The first few produce requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true.

Increase the number of retries in the producer configuration setting retries.

Performance degradation when SSL Is enabled

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with ssl.secure.random.implementation = SHA1PRNG. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

OPSAPS-43236: Kafka garbage collection logs are written to the process directory

By default Kafka garbage collection logs are written to the agent process directory. Changing the default path for these log files is currently unsupported.

None

RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transactional method because we are in an error state
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeFailWithError(TransactionManager.java:1125)
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeAddPartition(TransactionManager.java:442)
    at org.apache.kafka.clients.producer.KafkaProducer.doSend(KafkaProducer.java:1000)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:914)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:800)
    .
    .
    .
    Caused by: org.apache.kafka.common.errors.ClusterAuthorizationException: Cluster authorization failed.
```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting enable.idempotence to false.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

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

The Debezium Db2 Source connector does not support the evolution (updates) of schemas. In addition, schema change events are not emitted to the schema change topic if there is a change in the schema of a table that is in capture mode. For more information, see [DBZ-4990](#).

None.

CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy compression. The HDFS Stateless Sink connector is only affected if the Compression Codec or Compression Codec for Parquet properties are set to SNAPPY.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: 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 `libsnapy.so` was copied to the directory you created.
5. Remove the following from `/opt/nativelibs`.

```
libhadoop.a          libhadoop.so
                     libhadoop.so.1.0.0
```

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.

- Copy the following libraries from the downloaded archive to /opt/nativelibs on the cluster host.

```
libhadoop.a          libhadoop.so.1.0.0
```

The libraries are located in `hadoop-[***VERSION***]/lib/native`.

- Create a symlink named `libhadoop.so` and point it to `/opt/nativelibs/libhadoop.so.1.0.0`.

```
ln -s /opt/nativelibs/libhadoop.so.1.0.0 /opt/nativelibs/libhadoop.so
```

- Change the owner of every entry within `/opt/nativelibs` to `kafka`.

```
chown -h kafka:kafka /opt/nativelibs/*
```

- In Cloudera Manager, go to `Kafka service Configuration`.
- Add the following key-value pair to `Kafka Connect Environment Advanced Configuration Snippet (Safety Valve)`.
 - Key: `LD_LIBRARY_PATH`
 - Value: `/opt/nativelibs`
- Click `Save Changes`.
- Restart the `Kafka` service.

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

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to `required`. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to `requested` instead of `required` and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Unsupported Features

The following Kafka features are not supported in Cloudera Data Platform:

- Only Java and .Net based clients are supported. Clients developed with C, C++, Python, and other languages are currently not supported.
- The Kafka default authorizer is not supported. This includes setting ACLs and all related APIs, broker functionality, and command-line tools.
- SASL/SCRAM is only supported for delegation token based authentication. It is not supported as a standalone authentication mechanism.
- Kafka KRaft in this release of Cloudera Runtime is in technical preview and does not support the following:
 - Deployments with multiple log directories. This includes deployments that use JBOD for storage.
 - Delegation token based authentication.
 - Migrating an already running Kafka service from ZooKeeper to KRaft.
 - Atlas Integration.

Limitations

Collection of Partition Level Metrics May Cause Cloudera Manager's Performance to Degrade

If the Kafka service operates with a large number of partitions, collection of partition level metrics may cause Cloudera Manager's performance to degrade.

If you are observing performance degradation and your cluster is operating with a high number of partitions, you can choose to disable the collection of partition level metrics.



Important: If you are using SMM to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both SMM and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, SMM will not be able to display information about partitions. In addition, Cruise Control will not operate properly.

Complete the following steps to turn off the collection of partition level metrics:

1. Obtain the Kafka service name:
 - a. In Cloudera Manager, Select the Kafka service.
 - b. Select any available chart, and select Open in Chart Builder from the configuration icon drop-down.
 - c. Find \$SERVICENAME= near the top of the display.
The Kafka service name is the value of \$SERVICENAME.
2. Turn off the collection of partition level metrics:
 - a. Go to Hosts Hosts Configuration .
 - b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

Enter the following to turn off the collection of partition level metrics:

```
[KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_entity_update_enabled=false
```

Replace [KAFKA_SERVICE_NAME] with the service name of Kafka obtained in step 1. The service name should always be in lower case.

- c. Click Save Changes.

Known Issues in Knox

Learn about the known issues in Knox, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

CDPD-76435: Default HTTP client timeout for Streams Messaging Manager requests may cause connection timeouts

Some Streams Messaging Manager requests, particularly those retrieving topic data, may take up to 30 seconds to complete, even when there are no connectivity issues. Since the default gateway timeout is set to 20 seconds, this can result in connection timeouts despite the system functioning as expected.

Increase the HTTP client and socket timeouts for Streams Messaging Manager services to 60 seconds. This prevents connection timeouts when the requests take longer than the default gateway timeout.

Follow these steps to configure the required properties for the SMM API and SMM UI services:

1. In Cloudera Manager, select the Knox service.
2. Go to Configuration.

3. Find the Knox Simplified Topology Management - cdp-proxy property, and click **+** to add the following lines:

```
SMM-UI:replayBufferSize=512
SMM-UI:httpclient.connectionTimeout=1m
SMM-UI:client.socketTimeout=1m
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

4. Find the Knox Simplified Topology Management - cdp-proxy-api property, and click **+** to add the following lines:

```
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

5. Find the Knox Simplified Topology Management - cdp-proxy-token property, and click **+** to add the following lines:

```
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

6. Save your changes.
7. Refresh the cluster.

Known Issues in Kudu

Learn about the known issues in Kudu, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

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.

Workaround: Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

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.

Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

Unsupported features

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 in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

The following known issues were identified in this release:

No new known issues in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

No known issues identified in Livy.

Known issues in Navigator Encrypt

Learn about the known issues in Navigator Encrypt, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

Failed to start navencrypt-mount.service

After installing NavEncrypt on SLES, if the command "systemctl status navencrypt-mount" fails with the error: "Failed to start navencrypt-mount.service: Unit navencrypt-mount.service failed to load: No such file or directory", the systemd files need to be reloaded.

Run the command:

```
(sudo) systemctl daemon-reload
```

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

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 Oozie

Learn about the known issues in Oozie, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

There are no known issues identified in the release before Cloudera Runtime 7.1.9 SP1 CHF 11.

Known Issues in Ozone

Learn about the known issues in Ozone, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11**CDPD-87429: Unexpected log lines in Ozone snapshot diff output**

`ozone.client.server-defaults.validity.period.ms` configuration introduced to get server default configurations causes Replication Manager to fail when replicating data between Ozone buckets with the following error when the config value does not specify a timeunit.

```

` ` `
java.lang.IllegalStateException: Invalid first line in diff: "25
/07/02 05:58:44 INFO Configuration.deprecation: No unit for ozon
e.client.server-defaults.validity.period.ms(3600000) assuming MI
LLISECONDS"
  at com.cloudera.enterprise.distcp.OzoneCopyListingUtils.parseOz
oneDiff(OzoneCopyListingUtils.java:397)
` ` `

```

Overwrite the `ozone.client.server-defaults.validity.period.ms` to configure the value to have time unit (ns, ms, s, m, h). For example, 3600000ms.

CDPD-71648: Snapshot creation failing with "PERMISSION_DENIED" error on a bucket created during pre-upgrade in 7.1.7 SP3

On upgrade of cluster from 7.1.7 SP3 to 7.1.9 SP1, we are not able to create snapshots on a pre-upgrade volume or bucket with "PERMISSION_DENIED" error, since in 7.1.7SP3, we were using full KRB owner name with domain but currently, we only use and search for user shortname.

Update the owner of the bucket with the user shortname and then create snapshot by running the commands:

```

ozone sh bucket update --user=[***USER_SHORT_NAME***] {volume}/
{bucket}
ozone sh snapshot create {volume}/{bucket}
  {***SNAPSHOT_NAME***}

```

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

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

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

```

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

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

```

CDPD-64865: Intermittent timeout in TestBlockDeletion.testBlockDeletion (BlockDeletionService stuck)

If the event of SCM exiting safe mode was triggered multiple times, this might eventually cause the SCMBlockDeletingService to transition into the PAUSING state.

Check the SCM logs and ensure that the SCMBlockDeletingService is not running, and then restart the SCM.

CDPD-81569: Recon is not reflecting correct metrics for Volumes, Buckets, Keys, and Pending Key Deletions

Recon is displaying data for container and space utilisation only. The metrics for Volumes, Buckets, Keys, and Pending Key Deletions are not being updated and displayed correctly.

Apache Jira: [HDDS-12607](#)

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: 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.

The Prometheus binary is not available in CDP Private Cloud Base 7.1.9 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.

After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and if Ozone is in the Non-HA environment, an exception message is observed during the finalization of the Ozone upgrade.

During the finalization of the upgrade, a `ClassNotFoundException` message for `org.cloudera.log4j.redactor.RedactorAppender` class may be displayed. The error message can be ignored because the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

CDPD-68951: In CDP Private Cloud Base 7.1.9 CHF2 version and lower, the command `ozone sh key list <bucket_path>` displays the `isFile` flag in a key's metadata as `false` even when the key is a file. This issue is rectified in CDP Private Cloud Base 7.1.9 CHF3. However, the pre-existing (pre-upgrade) key's metadata cannot be changed.

None

CDPD-67771: When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
```

```

org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(OzoneManager
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.lambda
$1(OzoneManagerStateMachine.java:363) at
java.base/java.util.concurrent.CompletableFuture
$AsyncSupply.run(CompletableFuture.java:1700) at java.base/
java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1128)
at java.base/java.util.concurrent.ThreadPoolExecutor
$Worker.run(ThreadPoolExecutor.java:628) at java.base/
java.lang.Thread.run(Thread.java:834)

```

This occurs because S3A uses multipart upload to commit job results in a batch. The staging committer's replace mode deletes the target directory before completing Memory Protection Unit (MPU). The problem is that File System Optimization (FSO) does not create intermediate directories during MPU, it does only for regular file/dir/key requests.

Use `fs.s3a.committer.name=magic` for affected versions.

CDPD-64398: In SCM with non-HA configuration, Secret key manager does not initialise. Hence, the startup of OM and DataNode fails because it cannot get a secret key. This key is used when security is enabled and used in Block token and container token verification while communicating between the Ozone client, OM, and Datanode.

This issue does not occur in SCM with HA configuration.

You must force exit from the safe mode for SCM. This triggers the initialization of the Secret key manager. Below are the options:

- exit safe mode manually whenever SCM is started
- disable safe mode by setting `hdds.scm.safemode.enabled=false` in safety valve for SCM configuration

Impact of disabling safe mode and the purpose of safe mode:

- Write should not fail once the cluster is out of safe mode
- Read of existing data should not fail after the cluster is out of safe mode
- Unnecessary re-replication should be avoided during cluster restart

So disabling safe mode does not have any major impact on Ozone.

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.

CDPD-52412: keyManagerImpl#listStatus exceeds the maximum RPC length.

To reduce the payload, you must increase the part size.

OPSAPS-68159: If you did not deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Error when distributing parcel to host, Ignoring non-compliant parcel manifest error is displayed after Cloudera Manager is upgraded to 7.11.3.

If you encounter the error, perform the following steps:

1. Deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.11.3.

2. Restart the cluster with a delay of 10 minutes.
3. Continue to upgrade CDH 7.1.8 to CDH 7.1.9.

OPSAPS-68159: If you did not deactivate the Ozone parcel 718.2.x on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Ozone roles during the CDH 7.1.8 upgrade to CDH 7.1.9.

If you encounter the error, perform the following steps:

1. Deactivate the Ozone parcel 718.2.x.
2. Restart the Ozone service.
3. Perform Finalize Upgrade for the Ozone service.

Step result: The Ozone roles will be displayed in green.

CDPD-60989: The packaging version for Apache Ozone points to the 1.2.0 older version. This is a version string problem and not a packaging issue. The version of the Apache Ozone binary is closest to 1.4.0.

None. This only affects the JAR versioning.

CDPD-60366: Native library loader fails when system property native.lib.tmp.dir is not set. It fails because the library is copied to / instead of the cwd.

Setting native.lib.tmp.dir to a certain path like /tmp should solve the issue. To set this system property, you must add -Dnative.lib.tmp.dir=/tmp to ozone_java_opts in Cloudera Manager configuration of the Ozone cluster.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-60598: DataNode can consume up to the JVM heap size worth of direct memory buffers under specific failure scenarios. This can lead to a DataNode crash. This is due to Netty allocating and not freeing the direct memory buffers.

Restarting DataNode will resolve this issue.

CDPD-60466: The Ozone client is missing in the Cloudera Manager's Service Monitor node. This causes Ozone Canary's health check to fail.

Ensure you install Ozone CLI on the same node as Cloudera Manager's Service Monitor.

CDPD-60012: The Ozone SCM may be stuck in safe mode after upgrade to 7.1.9. This is due to a bug in how the SCM accounts for open storage containers.

The workaround is to temporarily set the Ozone configuration hdds.scm.safemode.threshold.pct to a lower value like 0.90 and restart the SCM.

CDPD-59679: In some scenarios, excess UNHEALTHY replicas of EC containers may not be removed from the cluster.

None

CDPD-50116: Topology-aware reads can provide better performance by routing applications to the closest replica of a file, if available, on the same physical node or same rack. However this feature is disabled by default in CDP 7.1.9.

This feature can be enabled by setting ozone.network.topology.aware.read to true using Cloudera Manager safety valve.

CDPD-57165: The DataNode disk usage thread may abort due to Ratis log tailing.

Use org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory to calculate disk usage. Configure in the ozone-site.xml property: hdds.datanode.du.factory value: org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory

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.

CDPD-55513: Hive external table replication policy does not migrate Hive tables in HDFS storage to Ozone.

To replicate data from HDFS to Ozone, see [Migrating your data from HDFS to Ozone](#). You can use HMS Mirror to replicate the metadata.

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"`

OPSAPS-67650: Ozone uses RocksDB as a library to persist metadata locally.

By default, RocksDB places certain executables in `/tmp`, and thus encounters errors when `/tmp` is mounted with `noexec`.

The workaround is to configure RocksDB to put executables at another location. On a PhatCat node, the steps are:

1. Go to Cloudera Manager UI OZONE Configuration .
2. Find Ozone Service Environment Advanced Configuration Snippet (Safety Valve) and set the following environment variable: `ROCKSDB_SHARED_LIB_DIR=/VAR/TMP`
3. Restart Ozone.

CDPD-49137: Ozone Manager Kerberos token expires for SCM communication and OM does not log in again.

Sometimes, OM's Kerberos token is not updated and it stops to communicate with SCM. When this occurs, writes start failing.

Restart OM or set the safety valve `hadoop.kerberos.keytab.login.autorenewal.enabled = true`

CDPD-56684: Keys get deleted when you do not have permission on volume

When a volume is deleted, it recursively deletes the buckets and keys inside it and only then deletes the volume. The volume delete ACL check is done only in the end, due to which you may end up deleting all the data inside the volume without having delete permission on the volume.



Note: There was no bucket/key permission available which allowed the user to delete them recursively.

CDPD-50610: Large file uploads are slow with OPEN and stream data approach

Hue file browser uses the append operation for large files. This API is not supported by Ozone in 7.1.9 and therefore large file uploads can be slow or timeout from the browser.

Use native Ozone client to upload large files instead of the Hue file browser.

OPSAPS-64097: Ozone service restart failed at SCM

Stopping SCM service using Cloudera Manager can sometimes timeout and need a retry. The Cloudera Manager API waits for 90 seconds which is not sufficient under certain circumstances.

Retry the shutdown using Cloudera Manager if the SCM still shows as running after a refresh of the service status.

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-filessystem-hadoop3.jar, it is likely due to the presence of the user hdfs on the node prior to CDP parcel activation. As a result, the Cloudera Manager agent skips the initialization related to HDFS, leading to the non-creation of the /var/lib/hadoop-hdfs directory.

Create the directory “/var/lib/hadoop-hdfs” on all nodes followed by the deactivation and activation of the CDP parcel (deactivate and activate the Ozone parcel instead, in case Ozone parcel is used).

CDPD-50447: When SCM High Availability is enabled, each of the SCM web UIs report the host of the web UI as the leader of HA, and the other two as followers. This gives wrong information

Correct output is available by running the `ozone admin scm roles --service-id=<ID>` command.

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](#).

OPSAPS-62327: In an Ozone cluster without any gateway roles, Ozone is unable to deploy client configurations and displays the ConfigGenException error.

You must add the Ozone gateway roles to the cluster.

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-36389: The configurations `datanodes.involved.max.percentage.per.iteration` and `size.moved.max.per.iteration` are meant to limit the max number of DataNodes that is involved and maximum size that can move in an iteration. This bug causes balancer to stop an iteration when it is two DataNodes or 1 Container size (5GB) away from hitting these limits. However, these DataNodes can again be considered for balancing in the next iteration. This means the cluster is balanced after enough iterations, albeit a bit slowly. This bug is apparent in small clusters of around four DataNodes where the DataNodes can be either the source or target for a lot of moves but the iteration gets stopped when three DataNodes have been involved. It takes a higher number of iterations to eventually balance this cluster. While this is a performance issue, it does not prevent balancer from ultimately balancing the cluster.

Increase the speed for balancing by decreasing the interval between each iteration using the configuration `balancing.iteration.interval`. The value of this configuration must be greater than `hdds.datanode.du.refresh.period`. `size.moved.max.per.iteration` can be increased to allow more data to move in one iteration.

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-35141: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2> (state=08S01,code=40000) java.sql.SQLException: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2>. This issue does not occur if the source and target buckets are different in Hive queries. Currently, copying across the same bucket is only supported.

Avoid different buckets in source and target path. Ensure that hive.exec.stagingdir is unset or set to the same bucket as the warehouse so that staging directories are created in the same bucket as the warehouse.

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-41184: With Legacy buckets, FileSystem Optimised is not interoperating with the Ozone shell command. The cause is that the directory key entry in the DB KeyTable is stored as "dir1/" with trailing slash. But while performing the described operation, Ozone shell (o3://) is normalizing the given path and removing the trailing slash "/" from it. That results in KEY_NOT_FOUND exception.

There are three workarounds:

- Use FileSystem API to Delete the Directories rather than Shell-Command API.
- Use FSO buckets instead of Legacy Buckets. As in FSO, you can create Intermediate Directories and Delete Directories using the Ozone shell commands.
- Disable and set the configuration ozone.om.enable.filesystem.pathsflag to false in order to delete the directories. This is generally not a preferred workaround because the cluster must be restarted to pick up the new changes.

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-12542: Mount of Ozone filesystem with the help of FUSE fails.

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-42897: EC writes are failing with "No enough datanodes to choose" after EC replication config is set globally.

EC writes start failing when large number of pipelines are created as a result of multiple EC configs with different chunk sizes used to write keys.

If standard EC configs (that is., rs-3-2-1024k) are used to write keys, number of pipelines created per DataNode is limited to five and this issue is not seen with standard EC configs.

The recommendation is not to create too many random chunk sizes. It is configurable because, users can decide based on their workload. But do not to have separate chunk sizes for each file.

CDPD-41539: The "No such file or directory" message is displayed when EC file is read using older of client.

You must upgrade the client before trying to read the key: `vol1/ecbuck1/1GB_ec`.

CDPD-40560: Filesystem Operations using Hadoop S3a connector on a FILE_SYSTEM_OPTIMIZED bucket fails. `org.apache.hadoop.ozone.om.exceptions.OMException: Unable to get file status: volume: s3v bucket: fso key: test/`

Do not run Hadoop S3a commands on a FILE_SYSTEM_OPTIMIZED bucket. Use OBJECT_STORE bucket layouts.

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.

CDPD-43432: Ozone Service in fault state in DataNode - Long Running setup.

Upgrade RocksDB to the latest version.

OPSAPS-63999: In the newly installed cluster, the Finish upgrade option is clickable.

None

OPSAPS-64648: Failed to start Ozone node using Cloudera Manager if default log path `/var/log/hadoop-ozone` does not exist. If this path does not exist, any Ozone nodes (for example SCM or DataNode) restart will fail.

Run the following command `sudo -u hdfs mkdir -p /var/log/hadoop-ozone` or replace `hdfs` with the user Ozone roles that are running.

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.

Known Issues in Spark

Learn about the known issues in Spark the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

The following known issues were identified in this release:

No new known issues in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the RuntimeWarning: Failed to add file message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

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

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

The following services support DDL statements:

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

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

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

CDPD-67517: Spark3 tests are failing if /tmp is mounted as NOEXEC.

Map the tmpdir to a writable path in spark3-conf/spark-defaults.conf using the following steps:

1. Go to the Cloudera Manager.
2. Go to the Spark service.
3. Click the Configuration tab.
4. Select Scope Gateway .
5. Select Category Advanced .
6. Locate the Spark Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-defaults.conf_client_config_safety_valve property.

7. Map the tmpdir to a writable path:

```
spark.driver.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
spark.executor.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
```

8. Enter a Reason for change, and then click Save Changes to commit the changes.
9. Deploy the client configuration.

CDPD-23817: In the upgraded Cluster, the permission of /tmp/spark is restricted due to the HDP configuration hive.exec.scratchdir=/tmp/spark.

If you are using the /tmp/spark directory in the CDP cluster, you must provide the required additional Policies/ACL permissions.

CDPD-22670 and CDPD-23103: There are two configurations in Spark, "Atlas dependency" and "spark_lineage_enabled", which are conflicted. The issue is when Atlas dependency is turned off but spark_lineage_enabled is turned on.

Run Spark application, Spark will log some error message and cannot continue. That can be restored by correcting the configurations and restarting Spark component with distributing client configurations.

CDPD-23007: Mismatch in the Spark Default DB Location. In HDP 3.1.5, hive_db entities have one attribute - 'location' which is configured to the '/managed' path. In fresh install of CDP 7.1.7, hive_db entities now have 2 attributes 'location' configured to '/external' path and 'managedLocation' configured to '/managed' path. In the AM2CM migration (HDP 3.1.5 -> CDP 7.1.7), the 'location' attribute from hive_db entities in HDP 3.1.5 comes unaltered to CDP 7.1.7 and hence maps to '/managed' path.

This issue arises only if you are upgrading from HDP 3.1.5 to CDP 7.1.7. If you are performing a fresh install of CDP 7.1.7, you can ignore this issue.

None

CDPD-217: The Apache Spark connector is not supported

The old *Apache Spark - Apache HBase Connector* (shc) is not supported in CDP releases.

Use the new HBase-Spark connector shipped in CDP release.

CDPD-3038: Launching pyspark displays several HiveConf warning messages

When pyspark starts, several Hive configuration warning messages are displayed, similar to the following:

```
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of name hive.meta
store.runworker.in does not exist
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of
name hive.masking.algo does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of
name hive.metastore.runworker.in does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of na
me hive.masking.algo does not exist
```

These errors can be safely ignored.

Known Issues in Schema Registry

Learn about the known issues in Schema Registry, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

CDPD-40380: Authorization checking issue when Kerberos is disabled

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the `ranger.plugin.schema-registry.service.name` field.
3. Replace `GENERATED_RANGER_SERVICE_NAME` with the actual name of the service.
4. Restart the Schema Registry service.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

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

Schema Registry uses a schema fingerprinting mechanism to differentiate between schemas. However, fingerprinting does not take into consideration the schema attributes of the field type. As a result, if you have two schemas where the only difference is that one has type attributes defined and the other does not, they will be considered identical by Schema Registry. For example, the following schemas are considered identical:

```
#Schema V1
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"local_timestamp_micros_long","type":"long"}]}

#Schema V2
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"local_timestamp_micros_long","type":{"type":"long","logicalType":"local-timestamp-micros"}}]}
```

Notice that the only difference is that in the second schema, the `local_timestamp_micros_long` field has a logical type specified. In cases like this, the new version of the schema is not created, the initial version is used. This is true even if the data that is being produced has a new schema version. The ID of the first schema version is used and is put in the serialized record. The new schema version is not created.

This issue is common when using change data capture (CDC) connectors like the Debezium connectors. This is because CDC connectors create schemas with the logical type decimal based on the column type in the database schema. For example:

```
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"database_column","type":{"type":"bytes","logicalType":"decimal","precision":64,"scale":0}}]}
```

If the database schema changes (for example, the column type), it is possible that only scale changes, which is a schema attribute.

```
{"type":"record","name":"schema_name","namespace":"ns","fields":[{"name":"database_column","type":{"type":"bytes","logicalType":"decimal","precision":64,"scale":1}}]}
```

In this case, even though scale changed to 1, the first version of the schema is used where scale is 0. As a result, the data is consumed with the wrong scale.

Avoid using logical types or other attributes. Alternatively, ensure that there are no changes in the logical types or other attributes between schema versions.

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

Schema Registry does not start if the address specified in the Load Balancer Address Ranger property does not end with a trailing slash (/).

Set the value of the `RANGER_REST_URL` Schema Registry environment variable to an address that includes a trailing slash.

1. In Cloudera Manager, select the Schema Registry service.
2. Go to Configuration.
3. Find the Schema Registry Server Environment Advanced Configuration Snippet (Safety Valve) property and add the following:

```
Key: RANGER_REST_URL
Value: [***RANGER REST API URL***]
```

Replace [***RANGER REST API URL***] with an address that can be used by Schema Registry to access Ranger. Ensure that the address ends with a trailing slash. For example: `http://ranger-1.cloudera.com:6182/`

4. Restart the Schema Registry service.

OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

1. Access the Ranger Console (Ranger Admin web UI).
2. Click the `cm_atlas` resource-based service.
3. Add the `schemaregistry` user to the all - * policies.
4. Click Manage Service Edit Service .
5. Add the `schemaregistry` user to the `default.policy.users` property.

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

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Known Issues in Sqoop

Learn about the known issues in Sqoop, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version

If you have upgraded to CDP PvC Base 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to CDP PvC Base 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see [Troubleshooting Apache Sqoop issues](#).

CDPD-44431: Using direct mode causes problems

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.

- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the `--direct` option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the `sqoop.enable.deprecated.direct` property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through `-Dsqoop.enable.deprecated.direct=true`.

CDPD-3089: Avro, S3, and HCat do not work together properly

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

Parquet columns inadvertently renamed

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.

PARQUET-99: Importing Parquet files might cause out-of-memory (OOM) errors

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.

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 in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager

Cloudera Manager does not display a Log Files menu for SMM UI role (and SMM UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by SMM UI is not supported by Cloudera Manager.

View the SMM UI logs on the host.

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in SMM safety valve in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Add the following value for bootstrap servers>Save Changes > Restart SMM.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

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

KNOX does not send neither the Authentication header nor the `hadoop.auth` cookie. Because of this, the SMM UI returns an HTTP 401 response and sets the `www-authenticate: Negotiate` header. After this, KNOX still does not send the Authentication header. This results in the SMM UI being inaccessible through Knox.

You can access the SMM UI directly in Cloudera Manager at `Clusters SMM Streams Messaging Manager Web UI` and log in using the Kerberos username and password.

CDPD-82560: Streams Messaging Manager known unprotected endpoint

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

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

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.snapshot` 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 in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

CDPD-11079: Blacklisted topics appear in the list of replicated topics

If a topic was originally replicated but was later excluded for replication, 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, this topic will be included in the response. Additionally, the excluded topic will also be visible in the SMM UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

CDPD-30275: SRM may automatically re-create deleted topics on target clusters

If `auto.create.topics.enable` is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with `srm-control`. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1]
```

2. Wait until SRM is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

Limitations

SRM cannot replicate Ranger authorization policies to or from Kafka clusters

Due to a limitation in the Kafka-Ranger plugin, SRM cannot replicate Ranger policies to or from clusters that are configured to use Ranger for authorization. If you are using SRM to replicate data to or from a cluster that uses Ranger, disable authorization policy synchronization in SRM. This can be achieved by clearing the Sync Topic Acls Enabled (`sync.topic.acls.enabled`) checkbox.

SRM cannot ensure the exactly-once semantics of transactional source topics

SRM data replication uses at-least-once guarantees, and as a result cannot ensure the exactly-once semantics (EOS) of transactional topics in the backup/target cluster.



Note: Even though EOS is not guaranteed, you can still replicate the data of a transactional source, but you must set `isolation.level` to `read_committed` for SRM's internal consumers. This can be done by adding `[***SOURCE CLUSTER ALIAS***]->[***TARGET CLUSTER ALIAS***].consumer.isolation.level=read_committed` to the Streams Replication Manager's Replication Configs SRM service property in Cloudera Manager. The `isolation.level` property can be set on a global connector or replication level. For example:

```
#Global connector level
connectors.consumer.isolation.level
=read_committed
#Replication level
uswest->useast.consumer.isolation.level
evel=read_committed
```

SRM checkpointing is not supported for transactional source topics

SRM does not correctly translate checkpoints (committed consumer group offsets) for transactional topics. Checkpointing assumes that the offset mapping function is always increasing, but with transactional source topics this is violated. Transactional topics have control messages in them, which take up an offset in the log, but they are never returned on the consumer API. This causes the mappings to decrease, causing issues in the checkpointing feature. As a result of this limitation, failover operations for transactional topics is not possible.

Known Issues in Ranger

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

UnsupportedClassVersionError

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 UnsupportedOperationException.

Remove the GraalJS library from the classpath.

CDPD-68806: The Revoke operation for users belonging to a group or role permission does not function as expected.

List command is listing all the tables even when the user permission is revoked. And also the command does not add any deny policy to Ranger for that specific user.

For example:

user_1 is the part of group_1

user_2 is the part of role_1

Create a Ranger policy for group_1 and role_1 to access all the resources in HBase.

Try to revoke the permission for user_1 and user_2 using HBase shell global revoke command:

```
revoke 'user_1'
```

```
revoke 'user_2'
```

But there is no deny policy added for that user. And access is still granted to user_1 and user_2.

This behavior is currently not supported in HBase shell. Must be handled manually using the Ranger policy change.

CDPD-68739: The revoke command does not work when using the HBase shell

While using the HBase shell, running the revoke command does not cancel the user permission. Users are able to perform actions even after running the revoke command.

For example:

When a user is granted permission to a particular table using the HBase grant command.

```
hbase:003:0> grant 'user_1', 'RWCA', 'some_abc_table'
```

And then execute the global revoke command for that user:

```
hbase:005:0> revoke 'user_1'
```

Permission for that user 'user_1' for the table 'some_abc_table' not got revoked.

None.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-58860: After upgrading CDP-7.1.8 to CDP-7.1.9, cdp-proxy token missing from Knox-Ranger policy.

As part of [OPSAPS-67480](#) in 719 default ranger policy is added from cdp-proxy-token topology, so that after a new installation of CDP-7.1.9, the Knox-Ranger policy includes cdp-proxy-token. However, upgrades do not add cdp-proxy-token to cm_knox policies automatically.

Manually add cdp-proxy-token to the Knox policy, using Ranger Admin Web UI.

1. Log in to Cloudera Manager Ranger Admin Web UI, as a Ranger administrator.

2. On Ranger Admin Web UI Service Manager Resource Knox , click cm_knox.
3. In Knox Policies , open the CDP Proxy UI, API and Token policy.
4. In Knox Topology*, add cdp-proxy-token.
5. Click Save.
6. Restart Ranger.

CDPD-58704: hadoop roll key / key delete command shows operation failed error when one KMS host is down, even when operation succeeds

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered kms instances. if one KMS instance is stopped (not deleted), client gets 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.

```
hbase:001:0> revoke 'hrt_11'
ERROR: org.apache.hadoop.hbase.coprocessor.CoprocessorException: HTTP 400 Error: processSecureRevokeRequest processing failed
    at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor.preRevoke(RangerAuthorizationCoprocesor.java:1309)
    at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor.preRevoke(RangerAuthorizationCoprocesor.java:1128)
    at org.apache.hadoop.hbase.master.MasterCoprocesorHost$162.call(MasterCoprocesorHost.java:1857)
    at org.apache.hadoop.hbase.master.MasterCoprocesorHost$162.call(MasterCoprocesorHost.java:1854)
    at org.apache.hadoop.hbase.coprocessor.CoprocessorHost$0bserverOperationWithoutResult.callObserver(CoprocessorHost.java:558)
    at org.apache.hadoop.hbase.coprocessor.CoprocessorHost.executeOperation(CoprocessorHost.java:631)
    at org.apache.hadoop.hbase.master.MasterCoprocesorHost.preRevoke(MasterCoprocesorHost.java:1854)
    at org.apache.hadoop.hbase.master.MasterRpcServices.revoke(MasterRpcServices.java:2740)
    at org.apache.hadoop.hbase.shaded.protobuf.generated.MasterProtos$MasterService$2.callBlockingMethod(MasterProtos.java)
    at org.apache.hadoop.hbase.ipc.RpcServer.call(RpcServer.java:387)
    at org.apache.hadoop.hbase.ipc.CallRunner.run(CallRunner.java:139)
    at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:369)
    at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:349)
```

None

CDPD-56741: Improvement in log message when jwtauth not used

The following exception is printed at startup only and it is not cluttering logs:

```
2023-05-30 06:18:40,127 ERROR org.apache.ranger.rms.security.RMSJwtAuthFilter:
quasar-pibgzl-1.quasar-pibgzl.root.hwx.site-startStop-1]:
Failed to initialize Ranger RMS JWT Auth Filter.
java.lang.Exception: RangerJwtAuthHandler:
```

```

Mandatory configs ('jwks.provider-url' & 'jwt.public-key')
are missing, must provide atleast one.
at org.apache.ranger.authz.handler.jwt.RangerJwtAuthHan
dler.initialize(RangerJwtAuthHandler.java:84)
~[ranger-authn-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
at org.apache.ranger.rms.security.RMSJwtAuthFilter.init
ialize(RMSJwtAuthFilter.java:77)
~[ranger-rms-common-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0
-186]
at jdk.internal.reflect.NativeMethodAccessorImpl.invoke
0(Native Method) ~[?:?]
at jdk.internal.reflect.NativeMethodAccessorImpl.invoke(N
ativeMethodAccessorImp

```

This exception is logged, since mandatory JWT auth filter configs are not provided (expected for some environments like y-cloud) Even though JWT auth filter is failed to initialised, it falls back to kerberos auth, hence no impact from auth perspective

NA

CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup

This is a warning log printed in catalina.out file when Ranger RMS server is initialised. Following exception is observed only in Oracle 19 setup. FileNotFoundException: /usr/share/java/oraclepki.jar

NA

CDPD-48975: Ranger KMS KTS to KMS DB migration : keys with the same name but different case are not migrated

KMS keys are not case sensitive

No workarounds. Such keys combination are very rare and migration doc was updated to check such keys before starting the migration.

CDPD-42598: Kafka policy creation allowed with incorrect permissions.

When creating a Kafka policy from the UI, the permissions "Idempotent write" and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource

None.

CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload:

```

{"resourceName": "classification", "userInput": "", "resources": {"classification": []}}

```

expectation is to return all the classifications . But the response has only "classification" Happens similarly for entity-label , entity-business-metadata.

None.

CDPD-40734: User allowed to insert data into a hive table when there is a deny policy on a table column.

A user is allowed to enter data into a table even if there is a deny policy present on one of the table columns.

Test scenario details:

```

Policy setup :-
policy 1 :- all access policy for hrt_qa, hive and impala
users
resources - database - * , table - * , column - *
users : hrt_qa, hive, impala

```

```

access - all access allowed
policy 2 :- policy on test_1.table_1 for hrt_5
users : hrt_5
resources : database - test_1, table - table_1, column - *
access :- all access allowed
policy 3 :- deny policy on test_1.table_1.c0 for hrt_5
users : hrt_5
resources : database - test_1, table - table_1, column -
c0
access - all access denied
data setup :-
database - test_1
table - table_1(c0 int, c1 int)

```

The user is able to insert data into the table.

None.

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 in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

COMPX-22434: Queue Manager fails to recreate Dynamic Queue Schedules (DQS) rule after exiting the mixed mode

QueueManager fails to recreate the DQS rule after exiting the Yarn mixed mode. This occurs when the cluster is upgraded from 7.1.9 GA+7.11.3 to 7.1.9 SP1 CHF9.

This error does not occur on a fresh install of 7.1.9 SP1 CHF 9, and also during 7.1.9 SP1 GA to 7.1.9 SP1 CHF9 upgrade.

None.

COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

None.

COMPX-4644: Queue capacity rounding problem when configuration is initially set via YARN

When setting the capacity scheduler configuration through the YARN/Cloudera Manager configuration, there may be capacity values that use multiple decimal places. This results in rounding/floating point precision discrepancies in the UI when trying to validate that all sibling capacities equal 100%. The UI looks like all the numbers add up to 100, but the validation still displays an error and does not allow to save the capacities. It is also observed that the capacity is being calculated as, for example, 99.9999999991 in the backend.

- Create queues within the UI, or
- Ensure that capacities configured through the Capacity Scheduler safety valve do not have more than one decimal place.

20202 Database migration after enabling opt-in migration

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to CDP 7.1.9, forcing migration from H2 to PostgreSQL database.
- Upgrade to CDP 7.1.9 CHF2, moving back to H2 database.
- Upgrade to CDP 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version CDP 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

CDPD-56559: MapReduce jobs can intermittently fail during a rolling upgrade.

During a rolling upgrade between CDP versions 7.1.8 and 7.1.9, MapReduce jobs may fail with message, `RuntimeException: native snappy library not available`. Although the native Snappy compression library is not loaded, a checkmark displays to indicate that the Snappy compression library is loading for NodeManagers that are pending upgrades. This causes the MapReduce jobs that are associated with the NodeManagers to fail. After the upgrade, the jobs work as expected. This issue only impacts rolling upgrades from before CDP 7.1.9 to a higher version.

None.

COMPX-12021 Queue Manager configurations on Scheduler Configuration page are not working

When setting the following properties on the YARN Queue Manager UI, the properties are set in the `capacity-scheduler.xml` file which does not have any effect on YARN. The properties need to be set in the `yarn-site.xml` file, which does not happen when you set them through YARN Queue Manager.

- "Maximum Application Priority" – `"yarn.cluster.max-application-priority"`
- "Enable Monitoring Policies" – `"yarn.resourcemanager.scheduler.monitor.enable"`
- "Monitoring Policies" – `"yarn.resourcemanager.scheduler.monitor.policies"`



Note: You can also set this property on the YARN configuration page in Cloudera Manager as "Capacity Scheduler Preemption".

- "Preemption: Observe Only" – `"yarn.resourcemanager.monitor.capacity.preemption.observe_only"`
- "Preemption: Monitoring Interval (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.monitoring_interval"`
- "Preemption: Maximum Wait Before Kill (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.max_wait_before_kill"`
- "Preemption: Total Resources Per Round" – `"yarn.resourcemanager.monitor.capacity.preemption.total_preemption_per_round"`
- "Preemption: Over Capacity Tolerance" – `"yarn.resourcemanager.monitor.capacity.preemption.max_ignored_over_capacity"`
- "Preemption: Maximum Termination Factor" – `"yarn.resourcemanager.monitor.capacity.preemption.natural_termination_factor"`
- "Enable Intra Queue Preemption" – `"yarn.resourcemanager.monitor.capacity.preemption.intra-queue-preemption.enabled"`

1. In Cloudera Manager, select the YARN service.
2. Click the Configuration tab.
3. Search for `yarn-site.xml`.
4. Under YARN Service Advanced Configuration Snippet (Safety Valve) for `yarn-site.xml`, add the corresponding parameter and value you need.
5. Click Save Changes.
6. Restart the YARN services.

COMPX-13177: QueueManager webapp requests fail with 'HTTP ERROR 400

`java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA'`

Products:

- Cloudera Manager for CDP Private

- Cloud Base Cloudera Manager for CDP Public Cloud

Context:

- Centos 7.8 and Redhat 7.8 operating systems, when FIPS support is enabled.

Problem:

- When attempting to display the Yarn Queue Manager interface, Cloudera Manager displays an error: "HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA".

1. Edit the file [/etc/default/cloudera-scm-server]
2. Around line 28, modify the line that starts with

```
#export CMF_OVERRIDE_TLS_CIPHERS=...
```

3. Remove the comment mark #.
4. Remove all ciphers with "3DES" in the name.
5. Save the file.
6. Restart the Cloudera Manager Server service.

Third-party applications do not launch if MapReduce framework path is not included in the client configuration

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the `mapreduce.application.framework.path` property is set to the appropriate value, but third-party applications with their own configurations does not launch.

Set the `mapreduce.application.framework.path` property to the appropriate configuration for third-party applications.

JobHistory URL mismatch after server relocation

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated the `mapred-site.xml` file that references the correct JobHistory Server.

YARN cannot start if Kerberos principal name is changed

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN does not start. In such cases, the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the `znode` and restart the YARN service.
- Use the `reset ZK ACLs` command. This also sets the `znodes` below `/rmstore/ZKRMStateRoot` to `world:anyone:cdwra` which is less secure.

Queue Manager does not open on using a custom user with a default Kerberos principal

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager `process_username` property matches the YARN `process_username` property.

Known Issues in Zeppelin

Learn about the known issues in Zeppelin the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

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.

BUG-125263: Zeppelin service move fails on clusters upgraded from HDP3.1.5

Resolve the circular symlink issue on the Zeppelin node by linking the conf directory to a new directory under /etc/zeppelin:

- # mkdir -p /etc/zeppelin/<version>/0
- # rm /usr/hdp/<version>/zeppelin/conf
- # ln -s /etc/zeppelin/<version>/0 /usr/hdp/<version>/zeppelin/conf

Where version is the HDP version. For example, 7.1.x-yyy. Restart the Zeppelin server in Ambari.

CDPD-3090: Due to a configuration typo, functionality involving notebook repositories does not work

Due to a missing closing brace, access to the notebook repositories API is blocked by default.

From the CDP Management Console, go to Cloudera Manager for the cluster running Zeppelin. On the Zeppelin configuration page (Zeppelin serviceConfiguration), enter shiro urls in the Search field, and then add the missing closing brace to the notebook-repositories URL, as follows:

```
/api/notebook-repositories/** = authc, roles[{{zeppelin_admin_group}}]
```

Click Save Changes, and restart the Zeppelin service.

CDPD-2406: Logout button does not work

Clicking the Logout button in the Zeppelin UI logs you out, but then immediately logs you back in using SSO.

Close the browser.

OPSAPS-59802: Zeppelin and Livy roles should be co-located on the same host.

When installing or upgrading to CDP Private Cloud Base, you must co-locate all Zeppelin and Livy roles on the same cluster host due to an issue with certificate generation.

Known Issues in ZooKeeper

Learn about the known issues in ZooKeeper the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 11.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 11

There are no new known issues identified in the previous releases.

Cumulative hotfix 10

Know more about the cumulative hotfix 10 for CDP 7.1.9 SP1. This cumulative hotfix was released on August 13, 2025.

Cloudera Runtime 7.1.9.1054-4 (Cumulative Hotfix 10) download URL:**Parcel Repository Location**

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1054/parcels/
```

What's new in Cloudera Runtime 7.1.9 SP1 CHF 10

Understand the functionalities and improvements to features of components in Cloudera Runtime 7.1.9 SP1 CHF 10.

Python 3.9 support for Impala on RHEL 8

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

Ranger mixed case group comparison

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

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

Upgraded Postgresql to 14.16

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

What's new in Kafka**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 and LDAP authentication

Two new Java options are introduced for Kafka brokers and Kafka Connect that enable you to specify an allow list of URLs for OAuth and LDAP authentication mechanisms. This gives you stricter control over which URLs can be accessed by Kafka for authentication. The following Java options are introduced:

- **OAuth** – `org.apache.kafka.sasl.oauthbearer.allowed.urls`
- **LDAP** – `com.cloudera.kafka.ldap.allowed.urls`

By default these options are empty, which means that a connection to any URL is allowed. The configuration properties you use in Cloudera Manager to set the options is different for Kafka brokers and Kafka Connect. Use the following configuration properties for Kafka brokers and Kafka Connect:

- For **Kafka brokers**, configure the options by adding them 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
```

```
-Dcom.cloudera.kafka.ldap.allowed.urls=HTTP://WWW.LDAP-EXAMPLE-1.COM,HTTP://WWW.LDAP-EXAMPLE-2.COM
```

- For **Kafka Connect**, configure the options by adding them 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.urls="HTTP://WWW.OAUTH-EXAMPLE-1.COM,HTTP://WWW.OAUTH-EXAMPLE-2.COM"
```

```
EXTRA_ARGS=-Dcom.cloudera.kafka.ldap.allowed.urls="HTTP://WWW.LDAP-EXAMPLE-1.COM,HTTP://WWW.LDAP-EXAMPLE-2.COM"
```

For more information, see:

- [Configure LDAP authentication for Kafka brokers](#)
- [Configuring OAuth2 authentication for Kafka brokers](#)
- [Configuring allowed LDAP and OAuth URLs](#)

What's new in Ozone

Ozone configuration recommendations for improved stability

Cloudera recommends you to update the following configurations to improve stability:

- ozone.scm.block.deletion.max.retry = 5 (default value is 1024)
- hdds.datanode.volume.min.free.space = 20GB (No default value)

Also, if hdds.datanode.volume.min.free.space.percent is configured, make sure that you have removed hdds.datanode.volume.min.free.space.percent from the configurations.

Fixed issues in 7.1.9 SP1 CHF 10

Know more about the cumulative hotfix 10 for 7.1.9 SP1.



Note: Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p1054.69605909.

CDPD-85295: Migrate commons-configuration to 2.x

The commons-configuration library is now upgraded in Zeppelin to address security issues and bugfixes.

CDPD-84529: Upgrade Commons BeanUtils to 1.11.0

The Commons BeanUtils library is now upgraded in Zeppelin to address security issues and bugfixes.

CDPD-82056: UI: when server response date fields as '0', UI shows as current time

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. Now, 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-88578: Resolved Missing Python Packages in Hue

Previously, Python packages were missing, leading to failures in Hue operations. This issue is now resolved, as the Hue environment now includes all necessary dependencies, eliminating the previously encountered errors.

CDPD-72557: Use Ozone Manager's KMS from client side when connecting to a cluster and dealing with encrypted data

Ozone client now uses server side Key Management Server (KMS) settings by default to encrypt or decrypt encryption keys with the cluster KMS server.

Apache JIRA: [HDDS-11227](#)

CDPD-84518: Expired secret key can abort leader OM startup

Previously, Ozone Manager (OM) start failure issues were caused by Null Pointer Exception (NPE) because of expired SCM managed secret keys. This issue is now resolved.

Apache JIRA: [HDDS-13234](#)

CDPD-66713: HDDS-9388 broke encryption

Previously, a file was not encrypted in an encrypted bucket caused by HDDS-9388 changes. This issue is now resolved.

Apache JIRA: [HDDS-10363](#)

CDPD-83223: Validate Container Balancer move timeout and replication timeout configurations

The `move.timeout - move.replication.timeout - event.timeout.datanode.offset` rule is now added 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

Previously, delete commands were getting discarded in the container balancer due to incorrect expiration time. With this fix, the container balancer can now send delete commands correctly.

Apache JIRA: [HDDS-13067](#)

CDPD-87933: DeleteBlocksCommandHandler thread stop for normal exception

Previously, block deletion was stopped due to invalid volume and container content, which resulted in continuous retry by the Storage Container Manager (SCM). This issue is now fixed.

Apache JIRA: [HDDS-13448](#)

CDPD-64642: DataNode Restart stuck in irrecoverable state with multiple exceptions

Previously, invalid container whose db path was missing in the memory container set, caused Null Pointer Exception for operations when accessing the container. This issue is now fixed.

Apache JIRA: [HDDS-10061](#)

CDPD-62073: OM Ratis Write - Move ACL check and Bucket resolution to preExecute

Previously, inconsistent behavior of Access Control List (ACL) validation on Ozone Manager (OM) nodes in high availability (HA) due to environmental issues could result in different results of data in those OMs and could have inconsistent data in OMs in HA. This issue is now fixed.

Apache JIRA: [HDDS-9388](#)

CDPD-77054: Exporting shell entities can cause the import to fail

Previously, when exporting Apache Hive tables created with Apache Spark, shell entities could be present in Apache Atlas. If these unresolved shell entities were included in the export zip file, the subsequent import operation failed.

This issue is now fixed. The import process now ensures that table entities are properly initialized and marked as complete, even if shell entities are present. This prevents failures during import when unresolved shell entities exist in the export.

CDPD-84502: Advanced Search is not working properly in the Atlas UI

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 the Advanced Search operations in the Atlas UI.

COMPX-21359: Capacity Scheduler UI queue filter is not working as expected

The Capacity Scheduler UI queue filter did not work as expected when submitting apps with leaf queue's name. This issue is now resolved and the Replication Manager now returns the full queue path regardless of whether the application was submitted with a leaf queue name or a full queue path.

Apache JIRA: [YARN-11662](#)

CDPD-83781: Upgraded Jetty version to 9.4.57.v20241219

The Jetty version is upgraded to 9.4.57.v20241219 due to CVE-2024-8184 and CVE-2024-13009.

CDPD-83429: 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.

CDPD-88100: Reduced Impala log verbosity for query details

Impala logs were excessively verbose, particularly for Remote Procedure Calls (RPCs) that fetched query runtime profiles and execution summaries.

This issue is now resolved by lowering the logging level for these RPCs to VLOG_QUERY (or VLOG(1))

Apache Jira: [IMPALA-13556](#)

CDPD-85268: Stable table renames during metadata invalidations

A system error previously occurred when renaming a table while also invalidating its metadata.

The issue is now resolved by fixing NullPointerException in Catalogd that may occur when renaming a table while also invalidating its metadata.

Apache Jira: [IMPALA-13631](#)

CDPD-83645: RSASSA-PSS signature algorithm support in kRPC

kRPC communication previously failed when certificates used the RSASSA-PSS signature algorithm.

This issue is now resolved by updating the code (through KUDU-3663) to correctly determine the hash algorithm for RSASSA-PSS signatures.

Apache Jira: [IMPALA-14038](#)

CDPD-83215: Consistent user identity with cookie authentication

When using Impala with proxy clients, such as Knox, and cookie-based authentication, user requests sometimes failed.

This issue is now resolved by ensuring that the Kerberos user identity is correctly set in the connection context for sessions authenticated through cookies.

Apache Jira: [IMPALA-14083](#)

CDPD-87887: Results for HWC direct read V2 mode

After upgrading to CDP 7.1.7, reading tables with Hive Warehouse Connector (HWC) in DIRECT_READ_V2 mode returned incorrect data, as deletes were not applied and updates showed duplicated rows.

This issue is now resolved by reverting a problematic change and applying a new correction. This restores the correct behavior of the DIRECT_READ_V2 mode, ensuring accurate data is returned after update or delete operations.

CDPD-83679: Stable ARRAY(*) function for mixed column types

Users previously encountered a system error when using the ARRAY(*) function on tables combining string and boolean column types. This occurred because the system failed to correctly determine a common data type for these mixed columns.

This issue is now resolved by ensuring that the system correctly identifies string as the common data type when processing mixed string and boolean columns. This prevents the previous error and allows the ARRAY(*) function to work as expected.

Apache Jira: [HIVE-28815](#)

CDPD-65287: Stable Hive client close operations

Hive JDBC clients sometimes failed when closing operations, especially after network timeouts or retries. This happened because a duplicate close request arrived after the operation was already completed.

This issue is now resolved by allowing the JDBC client to ignore duplicate close requests that arrive after an operation is already completed.

Apache Jira: [HIVE-25509](#)

CDPD-84215: DayOfWeek UDF results in consistent results irrespective of UTC / non-UTC timezones

The dayofweek() function previously returned inconsistent results for non-UTC timezones when the vectorized execution path was used.

This issue is now resolved by modifying the internal implementation of the dayofweek() user-defined function (UDF) for the Vectorized path execution in a way that now the date is not converted to UTC timezone before the UDF logic implementation.

Apache Jira: [HIVE-28075](#)

CDPD-84531: Connection starvation in HMS if datanucleus value generation fails

Database connections were sometimes not released as expected when internal processes failed.

This issue is now resolved to ensure continuous and reliable Hive Metastore operation by upgrading the datanucleus-rdbms component to version 5.2.13.

Apache Jira: [HIVE-28839](#)

CDPD-85086/CDPD-85497: Consistent property application for compaction jobs

When you set special rules for Hive compaction jobs, such as how to access logs, those rules weren't always applied to the part that updates statistics. This meant you could not see important logs for fixing problems, and your special settings were not fully used.

This issue is now resolved by making sure that all the special rules you set for compaction jobs are now also applied to the StatsUpdater job. As a result, you can easily check compaction logs for troubleshooting, and all your custom settings work as expected.

Apache Jira: [HIVE-28947](#)

CDPD-83683: String conversion for complex data types

Applying functions such as `md5()` to complex data types previously caused errors because these types could not be directly converted to strings.

This issue is resolved by allowing complex data types to be converted to strings through casting. Now, functions such as `md5()` will work correctly with your complex data.

Apache Jira: [HIVE-28825](#)

CDPD-85189: Open transactions from target cluster after Hive ACID table incremental replication job creates issues

After every incremental replication job run for the Hive ACID table replication policy, the transactions that remain open after the dump and load operation now get aborted on the target cluster.

This is applicable to only those transactions that were not in an open state during the dump operation on the source cluster.

Apache Jira: [HIVE-28975](#)

CDPD-84058: Commit message does not provide the required details

Previously, the commit event message did not provide details about whether the event was associated with the read or write operation. The commit text in the `notification_log` message is now enhanced, so that you can use the commit details during the dump process to decide whether you want to skip or include the commit during the next replication job run.

Apache Jira: [HIVE-28825](#)

CDPD-85723: IllegalStateException with Iceberg table with DELETE

Running a query on an Iceberg table fails with an `IllegalStateException` error in the following scenario:

- The Iceberg table has delete files for every data file (no data files without delete files) AND
- An anti-join operation is performed on the result of the Iceberg delete operation (IcebergDeleteNode or HashJoinNode)

This fix resolves the issue by setting the `TableRefIds` of the node corresponding to the Iceberg delete operation (``IcebergDeleteNode`` or ``HashJoinNode``) to only the table reference associated with the data files, excluding the delete files.

Apache Jira: [IMPALA-14154](#)

CDPD-88148: LEFT ANTI JOIN fails on Iceberg V2 tables with Delete files

Queries using a LEFT ANTI JOIN fail with an `AnalysisException` if the right-side table is an Iceberg V2 table containing delete files. For example, consider the following query:

```
SELECT * FROM table_a a
LEFT ANTI JOIN iceberg_v2_table b
ON a.id = b.id;
```

The error `Illegal column/field reference'b.input_file_name' of semi-/anti-joined table 'b' is displayed because semi-joined tuples need to be explicitly made visible for paths pointing inside them to be resolvable.`

The fix updates the `IcebergScanPlanner` to ensure that the tuple containing the virtual fields is made visible when it is semi-joined.

Apache Jira: [IMPALA-13888](#)

CDPD-87719: Ozone snapshot diff command is failing

The `snapdiff` operation could fail with a `Null Pointer Exception` 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-87776: Recon - Retrigger of build whole NSSummary tree task submission inconsistency

Executor may not run scheduled task or remove executor service:

```
org.apache.hadoop.ozone.recon.ReconUtils#triggerRebuild
```

Apache Jira: [HDDS-13428](#)

CDPD-59790: Fix LDB CLI for Ozone Snapshot DB

The Ozone Debug LDB CLI did not support any Ozone Snapshot databases. This issue is not resolved.

Apache Jira: [HDDS-9119](#)

CDPD-66744: Set Times API does not work with linked buckets

The setTimes API that operated on a linked bucket failed because of a bug. This issue is now resolved.

Apache Jira: [HDDS-10369](#)

CDPD-88172: Make ozone.snapshot.filtering.service.interval reconfigurable

ozone.snapshot.filtering.service.interval can be changed and reconfigured using the ozone admin reconfig CLI to enable or disable the SstFilteringService or change the interval for which the service is run.

Apache Jira: [HDDS-13464](#)

CDPD-88179: Exception handling for unchecked exception for deleteBlock command from SCM

This fix provides improved exception handling for unchecked exception for deleteBlock command from SCM.

Apache Jira: [HDDS-13451](#)

CDPD-87270: dt configuration dependency on secret key configuration

Adjusted the secret key expiry calculation (hdds.secret.key.expiry.duration = 9d) to account for delegation token lifetimes. This change ensures that tokens remain valid for their full configured duration, preventing premature key invalidation and improving stable authentication.

Apache Jira: [HDDS-13343](#)

CDPD-87378: RangerOzoneAuthorizer displays a Null Pointer Exception (NPE) when multiple instances of the classes are initialized

Previously, a NPE can be thrown during Ozone Manager re-initialization while Ozone Manager automatic bootstrap takes place. The follower Ozone Manager can fail to apply a transaction in such a case. This issue is now fixed.

Apache Jira: [HDDS-11472](#)

CDPD-87124: Upgraded the poi-ooxml version to 5.4.0+

The poi-ooxml Java library version is upgraded to 5.4.0 due to CVE-2025-31672.

CDPD-74753: Improvements for large scale deletion

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 deletion progress.

Apache JIRA: [HDDS-11605](#), [HDDS-11712](#), [HDDS-11713](#), [HDDS-11498](#), [HDDS-8888](#), [HDDS-8882](#), [HDDS-12087](#), [HDDS-11633](#), [HDDS-4368](#), [HDDS-11509](#), [HDDS-11711](#), [HDDS-11779](#), [HDDS-12238](#), [HDDS-11511](#), [HDDS-11512](#), [HDDS-12442](#), [HDDS-12474](#), [HDDS-11507](#), [HDDS-12320](#), [HDDS-12175](#), [HDDS-10407](#), [HDDS-10386](#), [HDDS-12975](#), [HDDS-10561](#), [HDDS-13103](#), [HDDS-13057](#), [HDDS-13259](#), [HDDS-13142](#), [HDDS-13079](#), [HDDS-13080](#), [HDDS-8869](#), [HDDS-10705](#), [HDDS-13489](#)

CDPD-84567: Backport upstream fixes for missing blocks issue

Backported the following upstream fixes:

Apache JIRA: [HDFS-15764](#), [HDFS-16985](#), [HDFS-17342](#)

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

CVE-2024-38827	VMware
CVE-2024-53990	Async HTTP Client
CVE-2024-36114	Aircompressor
CVE-2024-47561	Apache Avro
CVE-2023-39410	Apache Avro
CVE-2025-30065	Apache Parquet
CVE-2025-31672	Poi ooxml
CVE-2025-31651	Apache Tomcat
CVE-2025-48988	Apache Tomcat
CVE-2025-31650	Apache Tomcat
CVE-2025-27817	Apache Kafka
CVE-2025-27818	Apache Kafka
CVE-2025-27819	Apache Kafka

Behavioral changes in Cloudera Runtime 7.1.9 SP1 CHF 10

Know more about the behavioral changes in cumulative hotfix 10 for 7.1.9 SP1.

Impala

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](#)

Kafka

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.

Known issues in 7.1.9 SP1 CHF 10

The following section lists the known issues and limitations, the areas of impact, and workaround for each Runtime component.

Known Issues in Atlas

Learn about the known issues in Ozone the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

CDPD-93307: Hive import logs are truncated during bulk import

When running the `import-hive.sh` script to perform bulk Hive metadata imports, the logging process fails to retain the expected history. The `import-hive.log` file rotates prematurely when it reaches 10 MB and retains only one backup file (`import-hive.log.1`), instead of the expected 100 MB file size and 20 backup files.

Supply a custom `log4j` configuration file with the correct structure by using the `ATLAS_OPTS` environment variable before running the import script:

```
<?xml version="1.0" encoding="UTF-8" ?>
<!DOCTYPE log4j:configuration SYSTEM "log4j.dtd">
<log4j:configuration xmlns:log4j="http://jakarta.apache.org/lo
g4j/">
  <appender name="console" class="org.apache.log4j.ConsoleAppen
der">
    <param name="Target" value="System.out"/>
    <layout class="org.apache.log4j.PatternLayout">
      <param name="ConversionPattern" value="%d %-5p - [%t:
%x] ~ %m (%C{1}:%L)%n"/>
    </layout>
  </appender>
  <appender name="FILE" class="org.apache.log4j.RollingFileApp
ender">
    <param name="File" value="${atlas.log.dir}/${atlas.log.f
ile}"/>
    <param name="Append" value="true"/>
    <param name="maxFileSize" value="100MB"/>
    <param name="maxBackupIndex" value="20"/>
    <layout class="org.apache.log4j.PatternLayout">
      <param name="ConversionPattern" value="%d %-5p - [%t:
%x] ~ %m (%C{1}:%L)%n"/>
    </layout>
  </appender>
  <logger name="org.apache.atlas" additivity="false">
    <level value="info"/>
    <appender-ref ref="FILE"/>
  </logger>
  <logger name="org.apache.kafka.common.config.AbstractConfig"
additivity="false">
```

```

        <level value="error"/>
        <appender-ref ref="FILE"/>
    </logger>
    <root>
        <priority value="warn"/>
        <appender-ref ref="FILE"/>
    </root>
</log4j:configuration>

```

Export the ATLAS_OPTS variable to point to your custom file before running the import script:

```

export ATLAS_OPTS="-Dlog4j.configuration=file:/path/to/custom-hi
ve-import-log4j.xml"
./import-hive.sh

```

CDPD-70450: Impala SQL queries that include the “WITH” clause should populate lineage in Atlas

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-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-6565: Issue with ddlQueries in Atlas for Hive/Impala tables created in default/non-default db using DAS/HUE respectively

ddlQueries are not created for Impala origin table. ddlQueries created for Impala/Hive CTAS tables have difference in names as compared to same operation for Hive table. The Impala table has the query in the ddlQueries entity name.

None

CDPD-55301: ddlQueries and ALTERNATIVE_* lineage missing for Spark tables created through spark3-shell

The ddlQueries and ALTERNATIVE_* lineage missing for Spark tables created using spark3-shell.

None

CDPD-56085: 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

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*.

None

CDPD-58554: Discard audits of specific classification, label, and business metadata

Support to control audits for specific classification, label, and business metadata is not present in the Custom Audit Filters feature.

None

CDPD-40346: The ddlQueries and ALTERNATIVE_ADDCOLS lineage missing for Impala tables.

The ALTERNATIVE_ADDCOLS lineage has some issue when an Impala table is altered and the corresponding lineage is not created.

CDPD-67112: Import transforms do not work as expected when replacing a string which already has ":"

The character ":" is not supported in path replacements. The import succeeds but location remains unchanged. The character ":" must be avoided.

None

CDPD-69150: Unable to add labels or user defined properties in Japanese

Adding Japanese labels or user defined properties results in the error message: “Invalid label: ###, label should contain alphanumeric characters, _ or -”

None

CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties

Entities, such as `hdfs_path`, are not returned by the search when their indexable properties are searched using partial search terms made of Japanese or Chinese characters. Only exact matches return results when searching the indexed properties made of Japanese or Chinese characters.

None

CDPD-68191: Suggestions do not return the correct results when searching multiple Chinese characters

Free text search does not return results when searching Chinese phrases made of multiple characters. Partial searches return the correct results.

None.

CDPD-71219: Regression : Suggestions don't work for single character words on indexable attributes

When searching for entities whose name (entity names are indexable) is a single character, search results are returned but suggestions are not. When searching for entities whose description (entity descriptions are not indexable) is a single character, both search results and suggestions are returned.

None.

CDPD-67450: Table name renaming operation is not updating or creating iceberg_table entity

Renaming an Iceberg table does not update the corresponding Atlas entity.

None.

CDPD-67089: Export/Import: When a table with Ozone path is exported as "connected", only the Ozone key is exported.

When table with Ozone path is exported as "connected", only Ozone key is exported. Other Ozone entities, such as Ozone volume, Ozone bucket are not exported.

None.

CDPD-43772: Performance issues with Atlas service

If there are lot of update operations and the compression type of column families of `atlas_janus` table is SNAPPY, then the Kafka message processing might become slower.

- Consider setting compression type of column families of `atlas_janus` table as GZ.

OPSAPS-67783: During rolling upgrade one among two Atlas server failed to start but Cloudera Manager considered as success

Cloudera Manager marks the Execute command Start on service Atlas-1 as a success even when Atlas service had failed to start successfully. In such cases, Atlas logs give the exact reason for Atlas start-up failure.

None

CDPD-19358: "IsIndexable" and "isOptional" value of a typedef's attribute is modified post migration.

In HDP-265, `falcon_feed_creation` has "stored-in" attribute which has `IsIndexable` value set to True.

Post migration, "stored-in" is moved to `relationshipAttributeDefs` and has `IsIndexable` value set to False.

Similarly, in HDP-265 the table attribute of `hive_storage_desc` `isOptional` is set to true. Post migration, `isOptional` set to False.

None.

CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command

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

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

In converting content from Navigator to Atlas, the create time for Hive tables is not moved to Atlas.

None

CDPD-10574: Suggestion order doesn't match search weights

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

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) does not 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 the metadata collection from HMS. For example, if an ALTER TABLE operation in Spark that is changing a table name and is reporting 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 "@" does not 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, and use 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 does 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 logout 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 CDP.

None.

CDPD-1664: Guest users are redirected incorrectly

Authenticated users logging in to Atlas are redirected to the CDP 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: The 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, to add metadata tags to relationships that you wanted to make sure were unique in the system, you can design the relationship attribute with the property IsUnique equal true. However, in this release, the IsUnique attribute is not enforced.

None.

DOCS-13759: Tag Propagation stops after a certain depth while the lineage is being extended

When a tag is added to an entity at timestamp T1, the entities along the lineage to which the tag must be propagated is calculated at T1. Before tag propagation completes, if the lineage is extended, tag does not propagate to the entities in the extended lineage.

CDPD-41142: When a Kafka console consumer group is run, more than one update audits are seen

After running the console consumer with a consumer group, verify the consumer group entity created, along with the metrics and notifications for the consumer group and topic. The expected result can be: one ENTITY_CREATE audit and one ENTITY_UPDATE audit. But more than one ENTITY_UPDATE audits are seen.

CDPD-40165: Two audits are created for SPARK CTAS table

When following Spark queries are fired:

```
spark.sql("create table table1(id int)")
spark.sql("create table table2 as select * from table1")
```

HMS sends "ENTITY_CREATE" and "ENTITY_FULL_UPDATE_V2".

The extra ENTITY_FULL_UPDATE_V2 message received from HMS is sent as a part of ALTER_TABLE_ADDCOLS event from the HMS Hook side. This behaviour is observed only when the queries are run from Spark SQL and not when run the same queries are run from Beeline.

CDPD-39197: Debug metrics returns empty data

When debug metrics is enabled and some operations are performed, the response is empty

CDPD-36495: Updating legacyAttribute from False to True resets the initially created relationshipAttributes values

Creating types, entities, and to start, you must set the relationship with is_legacy_attribute value as **False**.

Later, update the value relationshipDef is_legacy_attribute to True.

For the entities that were created before updating the is_legacy_attribute to True, the relationship Attributes value is reset.

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-29409: Hive import: Suggestion suggests entity which is deleted.

Suggestions suggests tables of a database, which is a deleted entity.

CDPD-25152: Tag propagation through deferred actions consumes additional time as compared to default flow

The additional time might be due to the small overhead added to create / update task vertex and which is run in the background. This also depends on number of tasks queued to be executed.

CDPD-42954: Zeppelin notebook fails after enabling Atlas-HDFS hook

The Zeppelin notebooks are failing with errors after enabling Atlas-HDFS hook in the CDP cluster.

When the below properties are set for atlas-client.properties in Cloudera Manager:

- atlas.jaas.KafkaClient.option.keyTab
- atlas.jaas.KafkaClient.option.principal

Along with adding the properties in /etc/atlas/conf/atlas-application.properties, Cloudera Manager also adds these properties to atlas-application.properties for other services (like Spark).

Adding these properties interferes with the normal flow of the services (like Spark)

To enable HDFS lineage feature, instead of setting these properties through Cloudera Manager, users can manually add the properties directly in /etc/atlas/conf/atlas-application.properties

CDPD-10576: Deleted Business Metadata attributes appear in Search Suggestions

Atlas search suggestions continue to show Business Metadata attributes even if the attributes have been deleted.

None.

CDPD-39427: [HDFS Lineage]: When the input is a directory in case of put/copyfromLocal/cp/mv commands, lineage is not created even though the script succeeds.

When Source is a directory and target is a directory which is already present in Atlas, the command succeeds and inserts the data in the desired location, but lineage is not created.

DOCS-13760: System Attributes search, __classificationNames: Search with parent tag does not return entities associated to its children tags

System attribute search with __classificationNames = parent_tag returns entities associated to parent_tag only and not entities associated to its children tag.

Workaround: Instead of using system attribute, employ the basic search attribute "classification" which lists entities associated with inherited classifications.

CDPD-35818: Basic search with tag filter provides approximateCount as -1 when there is no match and is 0 otherwise

When the following search operations are performed:

- Faceted search with both tag filter and entity filter

The observed approximateCount is -1.

Here both entity filter and tag filter are present and when there is no match the response received is -1.

- Faceted search with only entity filter

Performing a basic search provides an approximate value of 0 when there is no match.

- Faceted search with only tag filter

Whenever there is tag filter in the query and there is no entity match, the approximateCount is -1 and if the tag filter is not available, the response approximateCount is 0

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export

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

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](#)

CDPD-76789: Creating tag with name description throws java.lang.ClassCastException

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

CDPD-70321: Atlas Parallel import is failing with various errors

During a parallel import-export activity with six iceberg table policies with exportOption as db1.* for all six exports, all import fail after the exports.

CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

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

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

CDPD-79099: Unable to import hive tables using import script in Public Cloud

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-88230: Bulk import operations currently do not construct lineage as expected

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

Bulk import operations do not create audits for the imported entities.

Known Issues in Hive

Learn about the known issues in Hive, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10**CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress**

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

CDPD-60770: Beeline Authentication Issue with Special Characters in Passwords

When LDAP is enabled, users cannot authenticate with Beeline if the password contains a special character. For example, the following string fails:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>;user=user@XXX;password='R3G#xpXyoy1M0Jb1'
```

Use the `-p` parameter to execute the Beeline command:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>; -n user@XXX -p 'R3G#xpXyoy1M0Jb1'
```

CDPD-43769: Inconsistent behavior with CHAR comparisons and string literals

In Hive, 'char' data to 'string' data comparison is considered as equal, but 'char' data to 'varchar' data comparison is considered as not equal if the data has trailing spaces. For example,

```
create table test(colchar char(10), colvarchar varchar(20), colstring string);
insert into test values ('a ', 'a ', 'a ');

select count(*) from test where colchar = colstring;
Output returns 1 because char and string data is considered as equal

select count(*) from test where colchar = colvarchar;
```

Output returns 0 because char and varchar data is considered as not equal

It is recommended that you use 'varchar or string' data type instead of 'char' data type for any character or string representation.

CDPD-42726: Wrong results for agg queries from stats after running multi-insert query on managed table

Multiple inserts at the same time into the same partition results in invalid stats. For example:

```
create table source(p int, key int, value string);
insert into source(p, key, value) values (101,42,'string42');

create table stats_part(key int,value string) partitioned by (p
int);
from source
insert into stats_part select key, value, p
insert into stats_part select key, value, p;

select count(*) from stats_part;
-- In this scenario, StatsOptimizer helps serving this query be
cause the result should be rowNum of the partition p=101.
-- The result is 1, however, the result should be 2.
```

Rewrite the multi-insert statement by merging the branches inserting into the same partition to one branch.

CDPD-57574: Query execution fails due to NullPointerException in DagUtils.setupQuickStart

Hive queries may fail at runtime with a NullPointerException while executing the Tez graph at DagUtils.setupQuickStart method. A part of the stack trace from the error is shown below:

```
ERROR : Failed to execute tez graph.
java.lang.NullPointerException: null at org.apache.hadoop.hive.
ql.exec.tez.DagUtils.setupQuickStart(DagUtils.java:1724)
```

The error usually occurs from an invalid plan that has cycles created by semi-join or map-join edges.

Disable some of the optimizations in order to remove the cycles from the plan. It is sufficient to set one of the following properties:

- set hive.tez.dynamic.partition.pruning=false;
- set hive.tez.dynamic.semijoin.reduction=false;
- set hive.auto.convert.join=false;

To avoid any undesired performance regressions, disable the optimization(s) only for specific queries and not globally (through Cloudera Manager) for the whole workload or cluster.

CDPD-43107: SemanticException for INSERT INTO statement when Hive Cost-based Optimizer (CBO) is disabled

If you are running the INSERT INTO query and hive.cbo.enable is set to "false", the query fails with a SemanticException. For example,

```
set hive.cbo.enable=false;

CREATE TABLE mytable (
id INT,
str STRING
);
```

```
INSERT INTO mytable (id, str) VALUES (1, 'a');
```

```
Output:
org.apache.hadoop.hive.ql.parse.SemanticException: 0:0 Expected
2 columns for insclause-0/default@mytable; select produces 1 col
umns. Error encountered near token 'a'
```

This issue occurs only when columns are specified for the target table in the query. INSERT INTO mytable values (1, 'a'); does not result in an exception.

Enable CBO and run the query.

CDPD-45370: Cartesian join fails with an empty table

When the configuration `hive.tez.cartesian-product.enabled` is set to true, query fails if one of the tables involved in a Cartesian join has zero rows. This results in an `IllegalArgumentException` with the message `Number of items is 0. Should be positive.`

To address this issue, set `hive.tez.cartesian-product.enabled` to false before running the query.

Apache Jira: [\[HIVE-26582\]](#)

CDPD-43957: HiveServer shuts down during replication due to high resource usage

During Hive replication (both bootstrap and incremental), you may notice that the HiveServer (HS2) shuts down periodically with the following error:

```
java.sql.SQLException: org.apache.hive.jdbc.ZooKeeperHiveClientE
xception: Unable to read HiveServer2 configs from ZooKeeper
  at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.ja
va:265)
  at org.apache.hive.jdbc.HiveDriver.connect(HiveDriver.java:107)
  at java.sql.DriverManager.getConnection(DriverManager.java:664)
  at java.sql.DriverManager.getConnection(DriverManager.java:247)
  at com.cloudera.enterprise.hive3qt.Hive3QueryTool$HiveOperation.
execute(Hive3QueryTool.java:682)
  at com.cloudera.enterprise.hive3qt.Hive3QueryTool.main(Hive3Qu
eryTool.java:935)
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException:
Unable to read HiveServer2 configs from ZooKeeper
  at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureCo
nnParams(ZooKeeperHiveClientHelper.java:177)
  at org.apache.hive.jdbc.Utills.configureConnParamsFromZooKeeper
(Utills.java:580)
  at org.apache.hive.jdbc.Utills.parseURL(Utills.java:391)
  at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.j
ava:263)
  ... 5 more
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: T
ried all existing HiveServer2 uris from ZooKeeper.
  at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.getServerHosts
(ZooKeeperHiveClientHelper.java:132)
  at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureCon
nParams(ZooKeeperHiveClientHelper.java:172)
  ... 8 mor
```

This issue occurs when you replicate at scale with an unbalanced cluster setup that has all the roles running on the Cloudera Manager host. As a result, Cloudera Manager ends up in a bottleneck situation because HS2 crashes, preventing further replication. The logs (`/var/log/messages`) indicate that there were a large number of Java processes running on the host, which exhausted the host's memory and triggered an Out Of Memory Killer process to stop HS2.

You can restart HS2 and reinitiate replication. However, the replication process may take a longer time to complete.

- Follow these recommendations when setting up your source and target clusters:
 - Ensure that no HS2, Hive metastore (HMS), DataNode, or NameNode roles are running on the Cloudera Manager host.
 - Ensure that you have multiple instances of HS2 and HMS roles on different nodes.
 - It is recommended that you have the HS2 and HMS roles on different nodes than the DataNodes or NameNodes.

These practices increase the possibility of Hive replication at scale completing successfully.

- Limit the scale of Hive replication by temporarily disabling replication policies.

CDPD-40730: Parquet change can cause incompatibility

Parquet files written by the parquet-mr library in CDP 7.1.9, where the schema contains a timestamp with no UTC conversion will not be compatible with older versions of Parquet readers. The effect is that the older versions will still consider these timestamps as they would require UTC conversions and will thus end up with a wrong result. You can encounter this problem only when you write Parquet-based tables using Hive, and tables have the non-default configuration `hive.parquet.write.int64.timestamp=true`.

None.

CDPD-41274: HWC + Oozie issue: Could not open client transport with JDBC Uri

Currently only Spark cluster mode is supported in the Oozie Spark Action with Hive Warehouse Connector (HWC).

Use Spark action in cluster mode.

```
<spark xmlns="uri:oozie:spark-action:1.0">
  ...
  <mode>cluster</mode>
  ...
</spark>
```

CDPD-26556 After an upgrade, querying a CTAS table under certain conditions might throw an exception

If you upgrade your Hive cluster from CDH 6 to CDP 7, create a CTAS table in the CDP cluster from a table you upgraded from CDH, you might see the following exception when you query the new table:

```
class org.apache.hadoop.io.IntWritable cannot be cast to class o
rg.apache.hadoop.hive.serde2.objectinspector.StandardUnionObject
Inspector$StandardUnion
```

This issue involves CDH-based tables having columns of complex types ARRAY, MAP, and STRUCT.

CDPD-23506: OutOfMemoryError in LLAP

Long running spark-shell applications can leave sessions in interactive Hiveserver2 until the Spark application finishes (user exists from spark-shell), causing memory pressure in case of a high number of queries in the same shell (1000+).

You must close spark-shell so that sessions are closed. Add the owner of the database or the tables as a user with read or read/write access to the tables directly.

CDPD-23041: DROP TABLE on a table having an index does not work

If you migrate a Hive table to CDP having an index, DROP TABLE does not drop the table. Hive no longer supports indexes ([HIVE-18448](#)). A foreign key constraint on the indexed table prevents dropping the table. Attempting to drop such a table results in the following error:

```
java.sql.BatchUpdateException: Cannot delete or update a parent
row: a foreign key constraint fails ("hive"."IDXS", CONSTRAINT "
```

```
IDXS_FK1" FOREIGN KEY ("ORIG_TBL_ID") REFERENCES "TBLS ("TBL_ID"
))
```

There are two workarounds:

- Drop the foreign key "IDXS_FK1" on the "IDXS" table within the metastore. You can also manually drop indexes, but do not cascade any drops because the IDXS table includes references to "TBLS".
- Launch an older version of Hive, such as Hive 2.3 that includes IDXS in the DDL, and then drop the indexes as described in [Language Manual Indexing](#).

Apache Issue: [Hive-24815](#)

CDPD-17766: Queries fail when using spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive.

Do not specify spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive remotely.

Workaround: specify principal=hive.server2.authentication.kerberos.principal as shown in the following syntax:

```
jdbc:hive://<host>:<port>/<dbName>;principal=hive.server2.authen
tication.kerberos.principal;<otherSessionConfs>?<hiveConfs>#<hive
eVars>
```

CDPD-13636: Hive job fails with OutOfMemory exception in the Azure DE cluster

Set the parameter hive.optimize.sort.dynamic.partition.threshold=0. Add this parameter in Cloudera Manager (Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml)

CDPD-10848: HiveServer Web UI displays incorrect data

If you enabled auto-TLS for TLS encryption, the HiveServer2 Web UI does not display the correct data in the following tables: Active Sessions, Open Queries, Last Max n Closed Queries

Known Issues in Hue

Learn about the known issues in Hue, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

CDPD-58978: Batch query execution using Hue fails with Kerberos error

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark queries in a batch mode launches a shell script Oozie job from Hue and this is not supported on a secure cluster.

There is no workaround. You can submit the queries individually.

OPSAPS-69659: Hue service fails on restart with "Unable to find pycogp2 2.5.4" error

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.

1. Log in to Cloudera Manager as an Administrator.

2. Go to Clusters Hue Configurations and add the following key and value in the Hue Service Environment Advanced Configuration Snippet (Safety Valve) field:

Key: PYTHONPATH

Value: `[***PYTHON-PATH***]`

Replace `[***PYTHON-PATH***]` with the actual location where you have installed Python. For example, `/opt/cloudera/parcels/CDH/lib/hue/build/env/lib/python3.8/site-packages`

3. Click Save Changes.
4. Restart the Hue service.

CDPD-59677: Unable to view Phoenix tables on the left assist in Hue

On clusters secured with Knox, you may not be able to see Phoenix tables on the left assist that are present under the default database (that is, an empty("") database).

None.

CDPD-58142: A query is not pre-populated in the Hue editor after clicking on the “Re Execute” button

When you click Re Execute to rerun a query from the Job Browser Queries Query Details page, the query does not get populated on the Hue editor, as expected.

None.

CDPD-54376: Clicking the home button on the File Browser page redirects to HDFS user directory

When you are previewing a file on any supported filesystem, such as S3 or ABFS, and you click on the Home button, you are redirected to the HDFS user home directory instead of the user home directory on the said filesystem.

None.

CDPD-41306: pip3.8 freeze command does not work and results into an error

You see the `/usr/lib/hue/build/env/bin/python: No such file or directory` error when you run the following command:

```
build/env/bin/pip3.8 freeze
```

Run the freeze command as follows by specifying the paths of python3.8 and pip3.8:

```
/opt/cloudera/parcels/CDH/lib/hue/build/env/bin/python3.8 /opt/c  
loudera/parcels/CDH/lib/hue/build/env/bin/pip3.8 freeze
```

CDPD-43293: Unable to import Impala table using Importer

Creating Impala tables using the Hue Importer may fail.

If you have both Hive and Impala services installed on your cluster, then you can import the table using by selecting the Hive dialect from Tables Sources .

If only Impala service is installed on your cluster, then go to Cloudera Manager Clusters Hue Configurations and add the following line in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini field:

```
[beeswax]  
max_number_of_sessions=1
```

DWX-8602: Unable to import a large CSV file from the local workstation

You may see an error message while importing a CSV file into Hue from your workstation, stating that you cannot import files of size more than 200 KB.

Upload the file to S3 or ABFS and then import them into Hue using the Importer.

CDPD-39330: Unable to use the pip command in CDP

You may not be able to use the `pip` command in CDP 7.1.7 or higher and may see the following error when using `pip` in a command: “ImportError: cannot import name `chardet`”.

Follow the steps listed on [Unable to use pip command in CDP](#).

CDPD-24294: Hue uses the unsafe-inline directive in its Content Security Policy (CSP) header

Hue 4 web interface uses the `unsafe-inline` directive in its CSP header. As a result, the application server does not set the CSP header in its HTTP responses, and therefore does not benefit from the additional protection against potential cross-site scripting issues and other modern application vulnerabilities which a properly configured CSP may provide. This could lead to application vulnerability.

Cloudera recommends deploying additional security measures such as a firewall within the Hue server to control allowed connections, and SSO-based authentications mechanisms such as LDAP or SAML.

OPSAPS-58927: Connection failed error when accessing the Search app (Solr) from Hue

If you are using Solr with Hue to generate interactive dashboards and for indexing data, and if you have deployed two Solr services on your cluster and selected the second one as a dependency for Hue, then Cloudera Manager assigns the hostname of the first Solr service and the port number of the second Solr service generating an incorrect Solr URL in the search section of the `hue.ini` file. As a result, you may see a “Connection failed” error when you try to access the Search app from the Hue web UI.

1. Log into Cloudera Manager as an Administrator.
2. Go to Clusters Hue service Configuration and add the following lines in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field:

```
[search]
# URL of the Solr Server
solr_url=http://[***HOSTNAME***]:[***PORT***]/solr/
```

For example:

```
solr_url=http://solr2:4567/solr/
```

3. Click Save Changes.
4. Restart the Hue service.

CDPD-41136: Importing files from the local workstation is disabled by default

Cloudera has disabled the functionality to import files from your local workstation into Hue because it may cause errors. You may not see the Local File option in the Type drop-down menu on the **Importer** page by default.

You can enable the functionality to import files from your local workstation by specifying the following parameter in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field in Hue configurations in Cloudera Manager:

```
[indexer]
enable_direct_upload=true
```

Invalid S3 URI error while accessing S3 bucket

The Hue Load Balancer merges the double slashes (`//`) in the S3 URI into a single slash (`/`) so that the URI prefix `"/filebrowser/view=S3A:/"` is changed to `"/filebrowser/view=S3A:/"`. This results in an error when you try to access the S3 buckets from the Hue File Browser through the port 8889.

The Hue web UI displays the following error: “Unknown error occurred”.

The Hue server logs record the “ValueError: Invalid S3 URI: S3A” error.

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.

You should be able to load the S3 browser from both 8888 and 8889 ports.

Alternatively, you can use the Hue server port 8888 instead of the load balancer port 8889 to resolve this issue.

Error while rerunning Oozie workflow

You may see an error such as the following while rerunning an already executed and finished Oozie workflow through the Hue web interface: E0504: App directory [hdfs://cdh/user/hue/oozie/workspaces/hue-oozie-1571929263.84] does not exist.

To resolve this issue, add the following property in the Hue Load Balancer Advanced Configuration Snippet:

1. Sign in to Cloudera Manager as an administrator.
2. Go to Clusters Hue service Configurations Load Balancer and search for the Load Balancer Advanced Configuration Snippet (Safety Valve) for httpd.conf field.
3. Specify MergeSlashes OFF in the Load Balancer Advanced Configuration Snippet (Safety Valve) for httpd.conf field.
4. Click Save Changes.
5. Restart the Hue Load Balancer.

CDPD-17956, CDPD-14167: Setting idle session timeout for Hue does not work when the cluster is secured using Knox SSO

If Hue is configured with desktop.auth.backend.KnoxSpnegoDjangoBackend as the Authentication Backend, then the automatic idle session logout that is set by configuring the `idle_session_timeout` property does not take effect. You may also see 404 error while accessing Hue from the Knox UI when the `idle_session_timeout` property is not set to -1.

None

Query history displays "Result Expired" message

You see the "Result Expired" message under the Query History column on the **Queries** tab for queries which were run back to back. This is a known behaviour.

None.

CDPD-43956: Manually replace UUID when importing Oozie workflows containing sub-workflows

When importing Oozie workflows that contain sub-workflows, you must replace all the UUID entries with unique new entries. If you change the UUID of a sub-workflow, you must update that reference in the parent workflow to avoid circular dependencies.

Hue limitation after upgrading from CDH to CDP Private Cloud Base

The `hive.server2.parallel.ops.in.session` configuration property changes from TRUE to FALSE after upgrading from CDH to CDP Private Cloud Base. Current versions of Hue are compatible with this property change; however, if you still would like to use an earlier version of Hue that was not compatible with this property being FALSE and shared a single JDBC connection to issue queries concurrently, the connection will no longer work after upgrading.

Hue UI is blank upon login after upgrading to CDP 7.1.7 from CDH 6

If your cluster was secured using Knox, and if you have upgraded from CDH 6 to CDP 7.1.7, then you may see a blank Hue screen. This could happen because the `knox_proxyhosts` parameter is newly introduced in CDP, and it is possible that this parameter is not configured in Cloudera Manager under Hue configuration.

Specify the host on which you have installed Knox in the Hue Knox Proxy Hosts configuration as follows:

1. Log in to Cloudera Manager as an Administrator.
2. Obtain the host name of the Knox Gateway by going to Clusters Knox service Instances .
3. Go to Clusters Hue service Configuration and search for the Knox Proxy Hosts field.



Note: Cloudera Manager displays the following warning if the Knox Proxy Hosts field is empty when Knox Gateway is enabled on the CDP cluster: The parameter `knox_proxyhosts` cannot be empty. This can happen if there are no Knox Gateways. Please set the `knox_proxyhosts` to the list of hosts that have Knox Gateways.

4. Specify the Knox Gateway hostname in the Knox Proxy Hosts field.
5. Click Save Changes and restart the Hue service.

Unsupported features

CDPD-59595: Spark SQL does not work with all Livy servers that are configured for High Availability

SparkSQL support in Hue with Livy servers in HA mode is not supported. Hue does not automatically connect to one of the Livy servers. You must specify the Livy server in the Hue Advanced Configuration Snippet as follows:

```
[desktop]
[spark]
livy_server_url=http(s)://[***LIVY-FOR-SPARK3-SERVER-HOST***]:
[***LIVY-FOR-SPARK3-SERVER-PORT***]
```

Moreover, you may see the following error in Hue when you submit a SparkSQL query: Expecting value: line 2 column 1 (char 1). This happens when the Livy server does not respond to the request from Hue.

Specify all different Livy servers in the `livy_server_url` property one at a time and use the one which does not cause the issue.

CDPD-18491: PySpark and SparkSQL are not supported with Livy in Hue

Hue does not support configuring and using PySpark and SparkSQL with Livy in CDP Private Cloud Base.

Importing and exporting Oozie workflows across clusters and between different CDH versions is not supported

You can export Oozie workflows, schedules, and bundles from Hue and import them only within the same cluster if the cluster is unchanged. You can migrate bundle and coordinator jobs with their workflows only if their arguments have not changed between the old and the new cluster. For example, hostnames, NameNode, Resource Manager names, YARN queue names, and all the other parameters defined in the `workflow.xml` and `job.properties` files.

Using the import-export feature to migrate data between clusters is not recommended. To migrate data between different versions of CDH, for example, from CDH 5 to CDP 7, you must take the dump of the Hue database on the old cluster, restore it on the new cluster, and set up the database in the new environment. Also, the authentication method on the old and the new cluster should be the same because the Oozie workflows are tied to a user ID, and the exact user ID needs to be present in the new environment so that when a user logs into Hue, they can access their respective workflows.



Note: Migrating Oozie workflows from HDP clusters is not supported.

Known Issues in Kafka

Learn about the known issues in Kafka, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker) in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Save Changes > Restart SMM.

The `offsets.topic.replication.factor` property must be less than or equal to the number of live brokers

The `offsets.topic.replication.factor` broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a `GROUP_COORDINATOR_NOT_AVAILABLE` error until the cluster size meets this replication factor requirement.

None

Requests fail when sending to a nonexistent topic with `auto.create.topics.enable` set to true

The first few produce requests fail when sending to a nonexistent topic with `auto.create.topics.enable` set to true.

Increase the number of retries in the producer configuration setting `retries`.

Performance degradation when SSL Is enabled

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with `ssl.secure.random.implementation = SHA1PRNG`. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

OPSAPS-43236: Kafka garbage collection logs are written to the process directory

By default Kafka garbage collection logs are written to the agent process directory. Changing the default path for these log files is currently unsupported.

None

RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transactional method because we are in an error state
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeFailWithError(TransactionManager.java:1125)
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeAddPartition(TransactionManager.java:442)
    at org.apache.kafka.clients.producer.KafkaProducer.doSend(KafkaProducer.java:1000)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:914)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:800)
    .
    .
    .
    Caused by: org.apache.kafka.common.errors.ClusterAuthorizationException: Cluster authorization failed.
```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting `enable.idempotence` to `false`.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

DBZ-4990: The Debezium Db2 Source connector does not support schema evolution

The Debezium Db2 Source connector does not support the evolution (updates) of schemas. In addition, schema change events are not emitted to the schema change topic if there is a change in the schema of a table that is in capture mode. For more information, see [DBZ-4990](#).

None.

CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy compression. The HDFS Stateless Sink connector is only affected if the `Compression Codec` or `Compression Codec` for Parquet properties are set to `SNAPPY`.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: org.apache.hadoop.util.NativeCodeLoader.buildSupportsSnappy()
```

```
Failed to write to HDFS due to java.lang.RuntimeException: native snappy library not available: this version of libhadoop was built without snappy support.
```

Download and deploy missing libraries.



Important: Ensure that you complete steps 1-11 on all Kafka Connect hosts. Additionally, ensure that the advanced configuration snippet in step 12 is configured for all Kafka Connect role instances.

1. Create the /opt/nativelibs directory.

```
mkdir /opt/nativelibs
```

2. Change the owner to kafka.

```
chown kafka:kafka /opt/nativelibs
```

3. Locate the directory containing the Hadoop native libraries and copy its contents to the directory you created.

```
cp /opt/cloudera/parcels/CDH/lib/hadoop/lib/native/* /opt/nativelibs
```

4. Verify that 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

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Unsupported Features

The following Kafka features are not supported in Cloudera Data Platform:

- Only Java and .Net based clients are supported. Clients developed with C, C++, Python, and other languages are currently not supported.
- The Kafka default authorizer is not supported. This includes setting ACLs and all related APIs, broker functionality, and command-line tools.
- SASL/SCRAM is only supported for delegation token based authentication. It is not supported as a standalone authentication mechanism.
- Kafka KRaft in this release of Cloudera Runtime is in technical preview and does not support the following:
 - Deployments with multiple log directories. This includes deployments that use JBOD for storage.
 - Delegation token based authentication.
 - Migrating an already running Kafka service from ZooKeeper to KRaft.
 - Atlas Integration.

Limitations

Collection of Partition Level Metrics May Cause Cloudera Manager's Performance to Degrade

If the Kafka service operates with a large number of partitions, collection of partition level metrics may cause Cloudera Manager's performance to degrade.

If you are observing performance degradation and your cluster is operating with a high number of partitions, you can choose to disable the collection of partition level metrics.



Important: If you are using SMM to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both SMM and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, SMM will not be able to display information about partitions. In addition, Cruise Control will not operate properly.

Complete the following steps to turn off the collection of partition level metrics:

1. Obtain the Kafka service name:
 - a. In Cloudera Manager, Select the Kafka service.
 - b. Select any available chart, and select Open in Chart Builder from the configuration icon drop-down.
 - c. Find \$SERVICENAME= near the top of the display.

The Kafka service name is the value of \$SERVICENAME.

2. Turn off the collection of partition level metrics:
 - a. Go to Hosts Hosts Configuration .
 - b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

Enter the following to turn off the collection of partition level metrics:

```
[KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_entity_update_enabled=false
```

Replace [KAFKA_SERVICE_NAME] with the service name of Kafka obtained in step 1. The service name should always be in lower case.

- c. Click Save Changes.

Known Issues in Knox

Learn about the known issues in Knox, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

CDPD-76435: Default HTTP client timeout for Streams Messaging Manager requests may cause connection timeouts

Some Streams Messaging Manager requests, particularly those retrieving topic data, may take up to 30 seconds to complete, even when there are no connectivity issues. Since the default gateway timeout is set to 20 seconds, this can result in connection timeouts despite the system functioning as expected.

Increase the HTTP client and socket timeouts for Streams Messaging Manager services to 60 seconds. This prevents connection timeouts when the requests take longer than the default gateway timeout.

Follow these steps to configure the required properties for the SMM API and SMM UI services:

1. In Cloudera Manager, select the Knox service.
2. Go to Configuration.
3. Find the Knox Simplified Topology Management - cdp-proxy property, and click **+** to add the following lines:

```
SMM-UI:replayBufferSize=512
SMM-UI:httpclient.connectionTimeout=1m
SMM-UI:client.socketTimeout=1m
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

4. Find the Knox Simplified Topology Management - cdp-proxy-api property, and click **+** to add the following lines:

```
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

5. Find the Knox Simplified Topology Management - cdp-proxy-token property, and click **+** to add the following lines:

```
SMM-API:replayBufferSize=512
SMM-API:httpclient.connectionTimeout=1m
SMM-API:client.socketTimeout=1m
```

6. Save your changes.
7. Refresh the cluster.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

There are no known issues identified in earlier releases.

Known Issues in Kudu

Learn about the known issues in Kudu, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

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.

Workaround: Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

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.

Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

Unsupported features

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 in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

The following known issues were identified in this release:

No new known issues in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

No known issues identified in Livy.

Known Issues in Iceberg

Learn about the known issues in Iceberg, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

There are no known issues identified in earlier releases.

Known Issues in Impala

Learn about the known issues in Impala, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10**CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress**

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the `--tls_ciphersuites` startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.

Ensure that the list in the `--tls_ciphersuites` startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

Apache Jira: [IMPALA-14625](#)

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work

If you try to run `impala-shell` on RHEL 9 by setting the default python executable available in `PATH` to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, `IMPALA_PYTHON_EXECUTABLE=python3`.

CDPD-81980: Impala query may fail due to backpressure

Impala queries could fail with a dropped due to backpressure error, even when service queue and memory usage appear normal.

Increase the value of the `--mem_limit` startup option.

CDPD-42958: After upgrading the CDH 7.1.9 from CDH 6.x, under certain conditions you cannot insert data into a table

Under the following conditions, after upgrading from CDH 6.x to CDH 7.1.9 you cannot insert data into a table from Impala:

- On CDH 6.x, you created a database with Impala in a user specified HDFS location.
- Using Hive, you then created a table in the database.

Under these conditions, the database and table are stored in the user-specified HDFS directory. After upgrading, the HDFS directory of the table is read-only for Impala. Consequently, from Impala you cannot insert new data into the table because Impala does not have write permission on the HDFS directory.

Workaround: To resolve this issue, use either one of the following workarounds:

- Using the Ranger Web UI, in the policy repository `cm_hdfs`, grant the user 'impala' write permission on the directory where the table resides.
- Enter the following command to grant write permission to user 'impala' on the HDFS directory where the table resides.

```
hdfs dfs -setfacl -m default:user:impala:rwX <HDFS directory>
```

CDPD-87697 : Python 3.11 Support in impala-shell

Python 3.11 is not supported in the CDP 7.1.9 SP1 CHF9 versions of the `impala-shell`

CDPD-80019: Impala fails to start due to large catalog-update topic

Impala coordinators fail to start if the catalog-update statestore topic exceeds 2GB, hitting the Thrift 0.16.0 message size limit (`thrift_rpc_max_message_size`). This prevents the coordinator from processing the update.

- Enable local catalog mode: `--use_local_catalog=true`
- In Cloudera Manager, go to Impala Configuration search for "Catalog Server Command Line

Argument Advanced Configuration Snippet (Safety Valve)" and click **+**

- Add property name and value:

```
Name : --catalog_topic_mode
Value: minimal
```

- Save changes.

Apache Jira: [IMPALA-12350](#)

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 HDF
S file hdfs://nameservice1/warehouse/tablespace/managed/hive/<da
tabase>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException:
File does not exist: /warehouse/tablespace/managed/hive/<data
base>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

You must use a single Knox node to access Impala UI.

Impala api calls via kinox require configuration if the kinox customized kerberos principal name is a default service user name

To access impala api calls via kinox, if the kinox customized kerberos principal name is a default service user name, then configure "authorized_proxy_user_config" by clicking Clusters->impala->configuration. Include the kinox customized kerberos principal name in the comma separated list of values <knox_custom_kerberos_principal_name>="*" where <knox_custom_kerberos_principal_name> is the value of the Kerberos Principal in the Knox service. Select Clusters>Knox>Configuration and search for Kerberos Principal to display this value.

CDPD-21828: Multiple permission assignment through grant is not working

None

Problem configuring masking on tables using Ranger

The following Knowledge Base article describes the behavior when we configure masking on tables using Ranger. This configuration works for Hive, but breaks queries in some scenarios for Impala.

For a workaround, see the following Knowledge Base article: [ERROR: "AnalysisException: No matching function with signature: mask\(FLOAT\)" when Impala jobs fail with the following error with signature: mask\(FLOAT\)](#)

IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXXXX"

IMPALA-532: Impala should tolerate bad locale settings

If the LC_* environment variables specify an unsupported locale, Impala does not start.

Add LC_ALL="C" to the environment settings for both the Impala daemon and the Statestore daemon.

Avro Scanner fails to parse some schemas

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

IMPALA-691: Process mem limit does not account for the JVM's memory usage

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the impalad daemon.

To monitor overall memory usage, use the `top` command, or add the memory figures in the Impala web UI `/memz` tab to JVM memory usage shown on the `/metrics` tab.

IMPALA-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Remove trailing semicolon from the Avro schema.

IMPALA-1652: Incorrect results with basic predicate on CHAR typed column

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the `RPAD()` function to blank-pad literals compared with `CHAR` columns to the expected length.

IMPALA-1821: Casting scenarios with invalid/inconsistent results

Using a `CAST()` function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

IMPALA-2005: A failed CTAS does not drop the table if the insert fails

If a CREATE TABLE AS SELECT operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

Drop the new table manually after a failed CREATE TABLE AS SELECT

IMPALA-2422: % escaping does not work correctly when occurs at the end in a LIKE clause

If the final character in the RHS argument of a LIKE operator is an escaped \% character, it does not match a % final character of the LHS argument.

None

IMPALA-2603: Crash: impala::Coordinator::ValidateCollectionSlots

A query could encounter a serious error if it includes multiple nested levels of INNER JOIN clauses involving subqueries.

None

IMPALA-3094: Incorrect result due to constant evaluation in query with outer join

An OUTER JOIN query could omit some expected result rows due to a constant such as FALSE in another join clause. For example:

```
explain SELECT 1 FROM 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;
+---+
|
```

[illegible]

IMPALA-3509: Breakpad minidumps can be very large when the thread count is high

The size of the breakpad minidump files grows linearly with the number of threads. By default, each thread adds 8 KB to the minidump size. Minidump files could consume significant disk space when the daemons have a high number of threads.

Add `-minidump_size_limit_hint_kb=size` to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then 2 KB per thread after that.) The minidump file can still grow larger than the "hinted" size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

IMPALA-4978: Impala requires FQDN from hostname command on Kerberized clusters

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the `gethostname()` system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a Kerberized cluster.

Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname`, only returns the short name, pass the command-line flag `##hostname=FULLY_QUALIFIED_DOMAIN_NAME` in the startup options of all Impala-related daemons.

IMPALA-6671: Metadata operations block read-only operations on unrelated tables

Metadata operations that change the state of a table, like `COMPUTE STATS` or `ALTER RECOVER PARTITIONS`, may delay metadata propagation of unrelated unloaded tables triggered by statements like `DESCRIBE` or `SELECT` queries.

None

IMPALA-7072: Impala does not support Heimdal Kerberos

None

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.

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.

Known issues in Navigator Encrypt

Learn about the known issues in Navigator Encrypt, the impact or changes to the functionality, and the workaround.

Failed to start navencrypt-mount.service

After installing NavEncrypt on SLES, if the command "systemctl status navencrypt-mount" fails with the error: "Failed to start navencrypt-mount.service: Unit navencrypt-mount.service failed to load: No such file or directory", the systemd files need to be reloaded.

Run the command:

```
(sudo) systemctl daemon-reload
```

Crash in Navigator Encrypt kernel module if --profile-file= option is used with ACLs

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 Ozone

Learn about the known issues in Ozone, the impact or changes to the functionality, and the workaround.

CDPD-87429: Unexpected log lines in Ozone snapshot diff output

``ozone.client.server-defaults.validity.period.ms`` configuration introduced to get server default configurations causes Replication Manager to fail when replicating data between Ozone buckets with the following error when the config value does not specify a timeunit.

```
...
java.lang.IllegalStateException: Invalid first line in diff: "25
/07/02 05:58:44 INFO Configuration.deprecation: No unit for ozon
e.client.server-defaults.validity.period.ms(3600000) assuming MI
LLISECONDS"
  at com.cloudera.enterprise.distcp.OzoneCopyListingUtils.parseOz
oneDiff(OzoneCopyListingUtils.java:397)
...
```

Overwrite the ``ozone.client.server-defaults.validity.period.ms`` to configure the value to have time unit (ns, ms, s, m, h). For example, 3600000ms.

CDPD-71648: Snapshot creation failing with "PERMISSION_DENIED" error on a bucket created during pre-upgrade in 7.1.7 SP3

On upgrade of cluster from 7.1.7 SP3 to 7.1.9 SP1, we are not able to create snapshots on a pre-upgrade volume or bucket with "PERMISSION_DENIED" error, since in 7.1.7SP3, we were using full KRB owner name with domain but currently, we only use and search for user shortname.

Update the owner of the bucket with the user shortname and then create snapshot by running the commands:

```
ozone sh bucket update --user=[***USER_SHORT_NAME***] {volume}/{bucket}
ozone sh snapshot create {volume}/{bucket}
  {***SNAPSHOT_NAME***}
```

OPSAPS-74668: ozone.snapshot.deep.cleaning.enabled and ozone.snapshot.ordered.deletion.enabled configurations are missing with CDP 7.1.9 SP1 CHF and Cloudera Manager 7.13.1

Two Ozone Manager configurations are missing while using CDP 7.1.9 SP1 CHF after upgrading Cloudera Manager version from 7.11.3 to 7.13.1.400.

If you are using CDP 7.1.9 SP1 CHF, before upgrading Cloudera Manager version from 7.11.3 to 7.13.1.400, add the following configs to Ozone Manager Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozone-site.xml so that Ozone Manager does not miss important config after the Cloudera Manager upgrade:

```
<property>
<name>ozone.snapshot.deep.cleaning.enabled</name>
<value>false</value>
</property>

<property>
<name>ozone.snapshot.ordered.deletion.enabled</name>
<value>true</value>
</property>
```

Known Issues in Ranger

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

UnsupportedClassVersionError

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.

CDPD-68806: The Revoke operation for users belonging to a group or role permission does not function as expected.

List command is listing all the tables even when the user permission is revoked. And also the command does not add any deny policy to Ranger for that specific user.

For example:

user_1 is the part of group_1

user_2 is the part of role_1

Create a Ranger policy for group_1 and role_1 to access all the resources in HBase.

Try to revoke the permission for user_1 and user_2 using HBase shell global revoke command:

```
revoke 'user_1'
```

```
revoke 'user_2'
```

But there is no deny policy added for that user. And access is still granted to user_1 and user_2.

This behavior is currently not supported in HBase shell. Must be handled manually using the Ranger policy change.

CDPD-68739: The revoke command does not work when using the HBase shell

While using the HBase shell, running the revoke command does not cancel the user permission. Users are able to perform actions even after running the revoke command.

For example:

When a user is granted permission to a particular table using the HBase grant command.

```
hbase:003:0> grant 'user_1', 'RWCA', 'some_abc_table'
```

And then execute the global revoke command for that user:

```
hbase:005:0> revoke 'user_1'
```

Permission for that user 'user_1' for the table 'some_abc_table' not got revoked.

None.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-58860: After upgrading CDP-7.1.8 to CDP-7.1.9, cdp-proxy token missing from Knox Ranger policy.

As part of [OPSAPS-67480](#) in 719 default ranger policy is added from cdp-proxy-token topology, so that after a new installation of CDP-7.1.9, the Knox Ranger policy includes cdp-proxy-token. However, upgrades do not add cdp-proxy-token to cm_knox policies automatically.

Manually add cdp-proxy-token to the Knox policy, using Ranger Admin Web UI.

1. Log in to Cloudera Manager Ranger Admin Web UI, as a Ranger administrator.
2. On Ranger Admin Web UI Service Manager Resource Knox, click cm_knox.
3. In Knox Policies, open the CDP Proxy UI, API and Token policy.
4. In Knox Topology*, add cdp-proxy-token.
5. Click Save.
6. Restart Ranger.

CDPD-58704: Hadoop roll key / key delete command shows operation failed error when one KMS host is down, even when operation succeeds

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered KMS instances. If one KMS instance is stopped (not deleted), client gets 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.

```
hbase:001:0> revoke 'hrt_11'
ERROR: org.apache.hadoop.hbase.coprocessor.CoprocessorException: HTTP 400 Error: processSecureRevokeRequest processing failed
    at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocessor.preRevoke(RangerAuthorizationCoprocessor.java:1309)
    at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocessor.preRevoke(RangerAuthorizationCoprocessor.java:1128)
    at org.apache.hadoop.hbase.master.MasterCoprocessorHost$162.call(MasterCoprocessorHost.java:1857)
    at org.apache.hadoop.hbase.master.MasterCoprocessorHost$162.call(MasterCoprocessorHost.java:1854)
    at org.apache.hadoop.hbase.coprocessor.CoprocessorHost$ObserverOperationWithoutResult.callObserver(CoprocessorHost.java:558)
```



```

    at org.apache.hadoop.hbase.coprocessor.CoprocessorHost.executeOperation(CoprocessorHost.java:631)
    at org.apache.hadoop.hbase.master.MasterCoprocessorHost.preRevoke(MasterCoprocessorHost.java:1854)
    at org.apache.hadoop.hbase.master.MasterRpcServices.revoke(MasterRpcServices.java:2740)
    at org.apache.hadoop.hbase.shaded.protobuf.generated.MasterProtos$MasterService$2.callBlockingMethod(MasterProtos.java)
    at org.apache.hadoop.hbase.ipc.RpcServer.call(RpcServer.java:387)
    at org.apache.hadoop.hbase.ipc.CallRunner.run(CallRunner.java:139)
    at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:369)
    at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:349)

```

None

CDPD-56741: Improvement in log message when jwtauth not used

The following exception is printed at startup only and it is not cluttering logs:

```

2023-05-30 06:18:40,127 ERROR org.apache.ranger.rms.security.RMSJwtAuthFilter:
    quasar-pibgzl-1.quasar-pibgzl.root.hwx.site-startStop-1]:
    Failed to initialize Ranger RMS JWT Auth Filter.
    java.lang.Exception: RangerJwtAuthHandler:
    Mandatory configs ('jwks.provider-url' & 'jwt.public-key')
    are missing, must provide atleast one.
    at org.apache.ranger.authz.handler.jwt.RangerJwtAuthHandler.initialize(RangerJwtAuthHandler.java:84)
    ~[ranger-authn-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
    at org.apache.ranger.rms.security.RMSJwtAuthFilter.initialize(RMSJwtAuthFilter.java:77)
    ~[ranger-rms-common-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
    at jdk.internal.reflect.NativeMethodAccessorImpl.invoke0(Native Method) ~[?:?]
    at jdk.internal.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl

```

This exception is logged, since mandatory JWT auth filter configs are not provided (expected for some environments like y-cloud) Even though JWT auth filter is failed to initialised, it falls back to kerberos auth, hence no impact from auth perspective

NA

CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup

This is a warning log printed in catalina.out file when Ranger RMS server is initialised. Following exception is observed only in Oracle 19 setup. FileNotFoundException: /usr/share/java/oraclepki.jar

NA

CDPD-48975: Ranger KMS KTS to KMS DB migration : keys with the same name but different case are not migrated

KMS keys are not case sensitive

No workarounds. Such keys combination are very rare and migration doc was updated to check such keys before starting the migration.

CDPD-42598: Kafka policy creation allowed with incorrect permissions.

When creating a Kafka policy from the UI, the permissions "Idempotent write" and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource

None.

CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload:

```
{"resourceName": "classification", "userInput": "", "resources": {"classification": []}}
```

expectation is to return all the classifications . But the response has only "classification" Happens similarly for entity-label , entity-business-metadata.

None.

CDPD-40734: User allowed to insert data into a hive table when there is a deny policy on a table column.

A user is allowed to enter data into a table even if there is a deny policy present on one of the table columns.

Test scenario details:

```
Policy setup :-
  policy 1 :- all access policy for hrt_qa, hive and impala
  users
    resources - database - * , table - * , column - *
    users : hrt_qa, hive, impala
    access - all access allowed
  policy 2 :- policy on test_1.table_1 for hrt_5
  users : hrt_5
  resources : database - test_1, table - table_1, column - *
  access :- all access allowed
  policy 3 :- deny policy on test_1.table_1.c0 for hrt_5
  users : hrt_5
  resources : database - test_1, table - table_1, column -
c0
  access - all access denied
data setup :-
  database - test_1
  table - table_1(c0 int, c1 int)
```

The user is able to insert data into the table.

None.

Known Issues in Schema Registry

Learn about the known issues in Schema Registry, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

CDPD-40380: Authorization checking issue when Kerberos is disabled

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the ranger.plugin.schema-registry.service.name field.
3. Replace GENERATED_RANGER_SERVICE_NAME with the actual name of the service.

4. Restart the Schema Registry service.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

CDPD-59015: Schema Registry does not create new versions of schemas even if the schema is changed

Schema Registry uses a schema fingerprinting mechanism to differentiate between schemas. However, fingerprinting does not take into consideration the schema attributes of the field type. As a result, if you have two schemas where the only difference is that one has type attributes defined and the other does not, they will be considered identical by Schema Registry. For example, the following schemas are considered identical:

```
#Schema V1
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [{"name": "local_timestamp_micros_long", "type": "long"}]}

#Schema V2
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [{"name": "local_timestamp_micros_long", "type": {"type": "long", "logicalType": "local-timestamp-micros"}}]}
```

Notice that the only difference is that in the second schema, the `local_timestamp_micros_long` field has a logical type specified. In cases like this, the new version of the schema is not created, the initial version is used. This is true even if the data that is being produced has a new schema version. The ID of the first schema version is used and is put in the serialized record. The new schema version is not created.

This issue is common when using change data capture (CDC) connectors like the Debezium connectors. This is because CDC connectors create schemas with the logical type decimal based on the column type in the database schema. For example:

```
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [{"name": "database_column", "type": {"type": "bytes", "logicalType": "decimal", "precision": 64, "scale": 0}}]}
```

If the database schema changes (for example, the column type), it is possible that only scale changes, which is a schema attribute.

```
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [{"name": "database_column", "type": {"type": "bytes", "logicalType": "decimal", "precision": 64, "scale": 1}}]}
```

In this case, even though scale changed to 1, the first version of the schema is used where scale is 0. As a result, the data is consumed with the wrong scale.

Avoid using logical types or other attributes. Alternatively, ensure that there are no changes in the logical types or other attributes between schema versions.

OPSAPS-68708: Schema Registry might fail to start if a load balancer address is specified in Ranger

Schema Registry does not start if the address specified in the Load Balancer Address Ranger property does not end with a trailing slash (/).

Set the value of the `RANGER_REST_URL` Schema Registry environment variable to an address that includes a trailing slash.

1. In Cloudera Manager, select the Schema Registry service.
2. Go to Configuration.

- Find the Schema Registry Server Environment Advanced Configuration Snippet (Safety Valve) property and add the following:

```
Key: RANGER_REST_URL
Value: [***RANGER REST API URL***]
```

Replace [***RANGER REST API URL***] with an address that can be used by Schema Registry to access Ranger. Ensure that the address ends with a trailing slash. For example: `http://ranger-1.cloudera.com:6182/`

- Restart the Schema Registry service.

OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

- Access the Ranger Console (Ranger Admin web UI).
- Click the `cm_atlas` resource-based service.
- Add the `schemaregistry` user to the `all - *` policies.
- Click `Manage Service Edit Service`.
- Add the `schemaregistry` user to the `default.policy.users` property.

OPSAPS-69317: Kafka Connect Rolling Restart Check fails if SSL Client authentication is required

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to `required`. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to `requested` instead of `required` and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Known Issues in Spark

Learn about the known issues in Spark the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

The following known issues were identified in this release:

No new known issues in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the `RuntimeWarning: Failed to add file` message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

CDPD-67517: Spark3 tests are failing if /tmp is mounted as NOEXEC.

Map the tmpdir to a writable path in spark3-conf/spark-defaults.conf using the following steps:

1. Go to the Cloudera Manager.
2. Go to the Spark service.
3. Click the Configuration tab.
4. Select Scope Gateway .
5. Select Category Advanced .
6. Locate the Spark Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-defaults.conf_client_config_safety_valve property.
7. Map the tmpdir to a writable path:

```
spark.driver.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
spark.executor.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
```

8. Enter a Reason for change, and then click Save Changes to commit the changes.
9. Deploy the client configuration.

CDPD-23817: In the upgraded Cluster, the permission of /tmp/spark is restricted due to the HDP configuration hive.exec.scratchdir=/tmp/spark.

If you are using the /tmp/spark directory in the CDP cluster, you must provide the required additional Policies/ACL permissions.

CDPD-22670 and CDPD-23103: There are two configurations in Spark, "Atlas dependency" and "spark_lineage_enabled", which are conflicted. The issue is when Atlas dependency is turned off but spark_lineage_enabled is turned on.

Run Spark application, Spark will log some error message and cannot continue. That can be restored by correcting the configurations and restarting Spark component with distributing client configurations.

CDPD-23007: Mismatch in the Spark Default DB Location. In HDP 3.1.5, hive_db entities have one attribute - 'location' which is configured to the '/managed' path. In fresh install of CDP 7.1.7, hive_db entities now have 2 attributes 'location' configured to '/external' path and 'managedLocation' configured to '/managed' path. In the AM2CM migration (HDP 3.1.5 -> CDP 7.1.7), the 'location' attribute from hive_db entities in HDP 3.1.5 comes unaltered to CDP 7.1.7 and hence maps to '/managed' path.

This issue arises only if you are upgrading from HDP 3.1.5 to CDP 7.1.7. If you are performing a fresh install of CDP 7.1.7, you can ignore this issue.

None

CDPD-217: The Apache Spark connector is not supported

The old *Apache Spark - Apache HBase Connector* (shc) is not supported in CDP releases.

Use the new HBase-Spark connector shipped in CDP release.

CDPD-3038: Launching pyspark displays several HiveConf warning messages

When pyspark starts, several Hive configuration warning messages are displayed, similar to the following:

```
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of name hive.meta
store.runworker.in does not exist
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of
name hive.masking.algo does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of
name hive.metastore.runworker.in does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of na
me hive.masking.algo does not exist
```

These errors can be safely ignored.

Known Issues in Sqoop

Learn about the known issues in Sqoop, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10**CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version**

If you have upgraded to CDP PvC Base 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to CDP PvC Base 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see [Troubleshooting Apache Sqoop issues](#).

CDPD-44431: Using direct mode causes problems

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the `--direct` option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the `sqoop.enable.deprecated.direct` property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through `-Dsqoop.enable.deprecated.direct=true`.

CDPD-3089: Avro, S3, and HCat do not work together properly

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

Parquet columns inadvertently renamed

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.

PARQUET-99: Importing Parquet files might cause out-of-memory (OOM) errors

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.

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 in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager

Cloudera Manager does not display a Log Files menu for SMM UI role (and SMM UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by SMM UI is not supported by Cloudera Manager.

View the SMM UI logs on the host.

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in SMM safety valve in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Add the following value for bootstrap servers>Save Changes > Restart SMM.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

CDPD-70313: KNOX does not send Authentication header on FIPS configuration

KNOX does not send neither the Authentication header nor the hadoop.auth cookie. Because of this, the SMM UI returns an HTTP 401 response and sets the www-authenticate: Negotiate header. After this, KNOX still does not send the Authentication header. This results in the SMM UI being inaccessible through Knox.

You can access the SMM UI directly in Cloudera Manager at Clusters SMM Streams Messaging Manager Web UI and log in using the Kerberos username and password.

CDPD-82560: Streams Messaging Manager known unprotected endpoint

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

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

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 in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

CDPD-11079: Blacklisted topics appear in the list of replicated topics

If a topic was originally replicated but was later excluded for replication, 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, this topic will be included in the response. Additionally, the excluded topic will also be visible in the SMM UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

CDPD-30275: SRM may automatically re-create deleted topics on target clusters

If auto.create.topics.enable is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with srm-control. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1]
```

2. Wait until SRM is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

Limitations

SRM cannot replicate Ranger authorization policies to or from Kafka clusters

Due to a limitation in the Kafka-Ranger plugin, SRM cannot replicate Ranger policies to or from clusters that are configured to use Ranger for authorization. If you are using SRM to replicate data

to or from a cluster that uses Ranger, disable authorization policy synchronization in SRM. This can be achieved by clearing the Sync Topic Acls Enabled (sync.topic.acls.enabled) checkbox.

SRM cannot ensure the exactly-once semantics of transactional source topics

SRM data replication uses at-least-once guarantees, and as a result cannot ensure the exactly-once semantics (EOS) of transactional topics in the backup/target cluster.



Note: Even though EOS is not guaranteed, you can still replicate the data of a transactional source, but you must set isolation.level to read_committed for SRM's internal consumers. This can be done by adding `[***SOURCE CLUSTER ALIAS***]->[***TARGET CLUSTER ALIAS***].consumer.isolation.level=read_committed` to the Streams Replication Manager's Replication Configs SRM service property in Cloudera Manager. The isolation.level property can be set on a global connector or replication level. For example:

```
#Global connector level
connectors.consumer.isolation.level
=read_committed
#Replication level
uswest->useast.consumer.isolation.l
evel=read_committed
```

SRM checkpointing is not supported for transactional source topics

SRM does not correctly translate checkpoints (committed consumer group offsets) for transactional topics. Checkpointing assumes that the offset mapping function is always increasing, but with transactional source topics this is violated. Transactional topics have control messages in them, which take up an offset in the log, but they are never returned on the consumer API. This causes the mappings to decrease, causing issues in the checkpointing feature. As a result of this limitation, failover operations for transactional topics is not possible.

Known Issues in YARN 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 in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

COMPX-22434: Queue Manager fails to recreate Dynamic Queue Schedules (DQS) rule after exiting the mixed mode

QueueManager fails to recreate the DQS rule after exiting the Yarn mixed mode. This occurs when the cluster is upgraded from 7.1.9 GA+7.11.3 to 7.1.9 SP1 CHF9.

This error does not occur on a fresh install of 7.1.9 SP1 CHF 9, and also during 7.1.9 SP1 GA to 7.1.9 SP1 CHF9 upgrade.

None.

COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

None.

COMPX-4644: Queue capacity rounding problem when configuration is initially set via YARN

When setting the capacity scheduler configuration through the YARN/Cloudera Manager configuration, there may be capacity values that use multiple decimal places. This results in

rounding/floating point precision discrepancies in the UI when trying to validate that all sibling capacities equal 100%. The UI looks like all the numbers add up to 100, but the validation still displays an error and does not allow to save the capacities. It is also observed that the capacity is being calculated as, for example, 99.9999999991 in the backend.

- Create queues within the UI, or
- Ensure that capacities configured through the Capacity Scheduler safety valve do not have more than one decimal place.

20202 Database migration after enabling opt-in migration

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to CDP 7.1.9, forcing migration from H2 to PostgreSQL database.
- Upgrade to CDP 7.1.9 CHF2, moving back to H2 database.
- Upgrade to CDP 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version CDP 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

CDPD-56559: MapReduce jobs can intermittently fail during a rolling upgrade.

During a rolling upgrade between CDP versions 7.1.8 and 7.1.9, MapReduce jobs may fail with message, `RuntimeException: native snappy library not available`. Although the native Snappy compression library is not loaded, a checkmark displays to indicate that the Snappy compression library is loading for NodeManagers that are pending upgrades. This causes the MapReduce jobs that are associated with the NodeManagers to fail. After the upgrade, the jobs work as expected. This issue only impacts rolling upgrades from before CDP 7.1.9 to a higher version.

None.

COMPX-12021 Queue Manager configurations on Scheduler Configuration page are not working

When setting the following properties on the YARN Queue Manager UI, the properties are set in the `capacity-scheduler.xml` file which does not have any effect on YARN. The properties need to be set in the `yarn-site.xml` file, which does not happen when you set them through YARN Queue Manager.

- "Maximum Application Priority" – `"yarn.cluster.max-application-priority"`
- "Enable Monitoring Policies" – `"yarn.resourcemanager.scheduler.monitor.enable"`
- "Monitoring Policies" – `"yarn.resourcemanager.scheduler.monitor.policies"`



Note: You can also set this property on the YARN configuration page in Cloudera Manager as "Capacity Scheduler Preemption".

- "Preemption: Observe Only" – `"yarn.resourcemanager.monitor.capacity.preemption.observe_only"`
- "Preemption: Monitoring Interval (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.monitoring_interval"`
- "Preemption: Maximum Wait Before Kill (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.max_wait_before_kill"`
- "Preemption: Total Resources Per Round" – `"yarn.resourcemanager.monitor.capacity.preemption.total_preemption_per_round"`
- "Preemption: Over Capacity Tolerance" – `"yarn.resourcemanager.monitor.capacity.preemption.max_ignored_over_capacity"`
- "Preemption: Maximum Termination Factor" – `"yarn.resourcemanager.monitor.capacity.preemption.natural_termination_factor"`
- "Enable Intra Queue Preemption" – `"yarn.resourcemanager.monitor.capacity.preemption.intra-queue-preemption.enabled"`

1. In Cloudera Manager, select the YARN service.
2. Click the Configuration tab.
3. Search for yarn-site.xml.
4. Under YARN Service Advanced Configuration Snippet (Safety Valve) for yarn-site.xml, add the corresponding parameter and value you need.
5. Click Save Changes.
6. Restart the YARN services.

COMPX-13177: QueueManager webapp requests fail with 'HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA'

Products:

- Cloudera Manager for CDP Private
- Cloud Base Cloudera Manager for CDP Public Cloud

Context:

- Centos 7.8 and Redhat 7.8 operating systems, when FIPS support is enabled.

Problem:

- When attempting to display the Yarn Queue Manager interface, Cloudera Manager displays an error: "HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA".

1. Edit the file [/etc/default/cloudera-scm-server]
2. Around line 28, modify the line that starts with

```
#export CMF_OVERRIDE_TLS_CIPHERS=...
```

3. Remove the comment mark #.
4. Remove all ciphers with "3DES" in the name.
5. Save the file.
6. Restart the Cloudera Manager Server service.

Third-party applications do not launch if MapReduce framework path is not included in the client configuration

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the mapreduce.application.framework.path property is set to the appropriate value, but third-party applications with their own configurations does not launch.

Set the mapreduce.application.framework.path property to the appropriate configuration for third-party applications.

JobHistory URL mismatch after server relocation

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated the mapred-site.xml file that references the correct JobHistory Server.

YARN cannot start if Kerberos principal name is changed

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN does not start. In such cases, the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the znode and restart the YARN service.

- Use the reset ZK ACLs command. This also sets the znodes below /rmstore/ZKRMStateRoot to world:anyone:cdw which is less secure.

Queue Manager does not open on using a custom user with a default Kerberos principal

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager process_username property matches the YARN process_username property.

Known Issues in Zeppelin

Learn about the known issues in Zeppelin the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

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.

BUG-125263: Zeppelin service move fails on clusters upgraded from HDP3.1.5

Resolve the circular symlink issue on the Zeppelin node by linking the conf directory to a new directory under /etc/zeppelin:

- # mkdir -p /etc/zeppelin/<version>/0
- # rm /usr/hdp/<version>/zeppelin/conf
- # ln -s /etc/zeppelin/<version>/0 /usr/hdp/<version>/zeppelin/conf

Where version is the HDP version. For example, 7.1.x-yyy. Restart the Zeppelin server in Ambari.

CDPD-3090: Due to a configuration typo, functionality involving notebook repositories does not work

Due to a missing closing brace, access to the notebook repositories API is blocked by default.

From the CDP Management Console, go to Cloudera Manager for the cluster running Zeppelin. On the Zeppelin configuration page (Zeppelin serviceConfiguration), enter shiro urls in the Search field, and then add the missing closing brace to the notebook-repositories URL, as follows:

```
/api/notebook-repositories/** = authc, roles[{{zeppelin_admin_group}}]
```

Click Save Changes, and restart the Zeppelin service.

CDPD-2406: Logout button does not work

Clicking the Logout button in the Zeppelin UI logs you out, but then immediately logs you back in using SSO.

Close the browser.

OPSAPS-59802: Zeppelin and Livy roles should be co-located on the same host.

When installing or upgrading to CDP Private Cloud Base, you must co-locate all Zeppelin and Livy roles on the same cluster host due to an issue with certificate generation.

Known Issues in ZooKeeper

Learn about the known issues in ZooKeeper the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 10.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 10

There are no new known issues identified in the previous releases.

Cumulative hotfix 9

Know more about the cumulative hotfix 9 for CDP 7.1.9 SP1. This cumulative hotfix was released on July 15, 2025.

Cloudera Runtime 7.1.9.1045 (Cumulative Hotfix 9) download URL:

Parcel Repository Location

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1045/parcels/
```

What's new in Cloudera Runtime 7.1.9 SP1 CHF 9

Understand the functionalities and improvements to features of components in Cloudera Runtime 7.1.9 SP1 CHF 9.

New Database Support

Postgres 17 is now supported starting from 7.1.9 SP1 CHF 9.

Rocky Linux support for Cloudera Runtime 7.1.9 SP1 CHF 9

Starting with the Cloudera Runtime 7.1.9 SP1 CHF 9 release, Cloudera Runtime 7.1.9 SP1 CHF 9 provides support for Rocky Linux. This update ensures seamless compatibility with Rocky Linux version 9.4, offering greater flexibility and platform options.

Impala now supports OAuth Authentication

OAuth authentication is now supported using OAuth JWT bearer tokens. For more information, see [OAuth Authentication](#)

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 7.13.1.400** or a higher 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](#).

Fixed issues in 7.1.9 SP1 CHF 9

Know more about the cumulative hotfix 9 for 7.1.9 SP1.



Note: Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p1045.67903105.

CDPD-83120: Solr authentication bypass issue

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](#)

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.

CDPD-82671: Log files for RocksDB ops should be cleaned more frequently for heavy load

The RocksDB log files are quite heavy when a lot of snapshot operations are involved. This fix cleans log files of RocksDB operations frequently for heavy loads.

CDPD-81939: Volume scanner should fail volume if rocksDB is inaccessible

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-81528: Cleanup expired incomplete MPUs

Any parts of a multipart-upload which are not committed are retained for a duration of 30 days by default. With this fix, MultipartUploadCleanupService cleans up the expired parts by aborting the upload. The following configurations can be used to tune the cleanup:

- ozone.om.open.mpu.cleanup.service.interval (default value=24h)
- ozone.om.open.mpu.cleanup.service.timeout (default value=300s)
- ozone.om.open.mpu.expire.threshold (default value=30d)
- ozone.om.open.mpu.parts.cleanup.limit.per.task (default value=1000)

Apache Jira: [HDDS-9095](#)

CDPD-66534: Support cross realm Kerberos out of box

ozone.om.kerberos.principal.pattern was not configured previously which prevented cross-cluster communication in Kerberized environment. This issue is fixed.

Apache Jira: [HDDS-10328](#)

CDPD-81786: Auto-compact tables which can tend to be large in size at intervals

The following configuration properties are added:

Configuration property	Default value	Description
ozone.compaction.service.enabled	false	Enable or disable a background job that periodically compacts rocksdb tables flagged for compaction.
ozone.om.compaction.service.run.interval	6h	A background job that periodically compacts rocksdb tables flagged for compaction. Units can be defined with postfix (ns,ms,s,m,h,d).
ozone.om.compaction.service.timeout	10m	A timeout value of compaction service. If this is set greater than 0, the service will stop waiting for compaction completion after this time. Units can be defined with postfix (ns,ms,s,m,h,d).

Configuration property	Default value	Description
ozone.om.compaction.service.columnfamilyTable	anyTable, fileTable, directoryTable, deleteTable, appendTable, directoryTable, multipartInfoTable	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-81764: Patch to replace the long strings set in spark_process attributes

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.

CDPD-80588: Tomcat Upgrade to 9.0.99+

Upgraded Tomcat to 9.0.99+ to address CVE-2025-24813.

CDPD-79160: Nullpointer exception while deleting business metadata

After the update, the migration status is stored correctly and reused to restart migration from the point where it failed earlier.

Apache Jira: [ATLAS-4863](#)

CDPD-83356: Improved database type detection for metastore

Hive Metastore clients sometimes experienced delays when connecting to the metastore.

This issue was addressed by enhancing how the system determines the database type during metastore initialization. This improvement avoids repetitive checks, reduces connection delays, and ensures HMS clients connect more quickly and reliably.

Apache Jira: [HIVE-28460](#)

CDPD-83355/CDPD-73669: Secondary pool connection starvation caused by updatePartitionColumnStatisticsInBatch API

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-81122: Enhanced concurrent access in HWC secure mode

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-82887: Hive web interfaces no longer expose server version

Hive web interfaces and related services previously exposed their underlying server version in the header.

This issue was addressed by stopping the web interfaces from sending information about their server version.

CDPD-83079: Included Iceberg fixes related to EXPIRE SNAPSHOTS

This release includes the following Iceberg fixes to ensure that the EXPIRE SNAPSHOTS can effectively delete data files from storage:

- Deletion of expired snapshot files in a transaction — When a snapshot is expired as part of a transaction, the manifest list files should be deleted upon transaction commit. Previously, a recent change prevented these files from being deleted if they were also committed as part of a transaction. However, this caused issues in simpler cases where no new files were committed.

This issue is now fixed by ensuring deletion is not skipped when the list of committed files is empty.

- Improved logic for determining committed files in Base Transaction — The logic to identify the set of committed files in 'BaseTransaction' has been corrected for scenarios where no new snapshots are available.

CDPD-83116: Query planning incorrectly estimated row counts

Impala's query planner generated negative cardinality estimation over very specific condition (TOP-N sort over very high NDV column). This results in `IllegalStateException: null` thrown and fail the query.

This issue was resolved by improving the calculation that estimates row counts for sorting operations.

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-67359: Improved S3 Directory Listing in Hue for Large Datasets

Previously, Hue encountered "incorrect signature" errors when attempting to list S3 directories containing 1,000 or more objects. This issue was caused by S3's pagination mechanism, where the marker field for subsequent requests was not correctly passed to the Ranger Authorization Service (RAZ). As a result, the signature generated by RAZ did not align with the actual S3 request, leading to verification failures.

This problem is fixed by enhancing Hue's handling of the S3 marker field. Hue now accurately includes this pagination parameter when requesting signed headers from RAZ. This ensures that RAZ provides a precise signature, enabling S3 to successfully verify the request and allow Hue to correctly list and paginate through large directories.

CDPD-82364: Nested loop join rewrites disjunctive subquery incorrectly

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-82303: EXEC_TIME_LIMIT_S incorrectly includes planning time

The `EXEC_TIME_LIMIT_S` setting was triggered during the planning and scheduling phases, which could cause queries to fail before any processing on backends started.

The issue was addressed by starting the `EXEC_TIME_LIMIT_S` countdown only after the query is ready to run on the backends. This ensures the timeout applies only to the actual processing phase.

Apache Jira: [IMPALA-14001](#)

CDPD-82251: Inconsistent row count output in Impala-shell

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-81556: Improved handling of invalid data in text files

Impala encountered crashes when processing text files that contained invalid binary data.

This issue was resolved by enhancing Impala's data handling. The system now correctly identifies and flags invalid binary data encountered in text files, preventing system instability.

Apache Jira: [IMPALA-13927](#)

CDPD-80019: Impala now starts with large catalog-update topics

Impala coordinators previously failed to start if their internal update data exceeded 2GB. This prevented the coordinator from processing necessary updates.

This issue was resolved.

Apache Jira: [IMPALA-13020](#)

CDPD-66938: [Analyze] [Atlas] test_time_range tests fail

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.

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

- [CVE-2025-22228](#) - Spring Security

Behavioral changes in Cloudera Runtime 7.1.9 SP1 CHF 9

Know more about the behavioral changes in cumulative hotfix 9 for 7.1.9 SP1.

Atlas

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 a large amount of data processed at Atlas.

To keep the data for `details` and `sparkPlanDescription`, `atlas.notification.consumer.preprocess.spark_process.attributes`, set them to false explicitly.

Hive

Summary:

Metastore secondary connection pool size is now configurable

Previous Behavior: The Metastore's secondary connection pool had a fixed size of 2. This often led to connection limitations, especially under heavy workloads.

New Behavior: You can now configure the metastore's secondary connection pool size using the property `datanucleus.connectionPool.secondary.maxPoolSize`. This lets you adjust the pool beyond its default of 2, preventing connection limitations and improving performance.

Apache Jira: <https://issues.apache.org/jira/browse/HIVE-28900>

Impala

Summary:

OAuth Authentication now supported

Previous Behavior: Authentication using OAuth was not available.

New Behavior: OAuth authentication is now supported using OAuth JWT bearer tokens. For more information, see [OAuth Authentication](#)

Known issues in 7.1.9 SP1 CHF 9

The following section lists the known issues and limitations, the areas of impact, and workaround for each Runtime component.

Known Issues in Atlas

Learn about the known issues in Ozone the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 9.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9

CDPD-77054: Exporting shell entities can cause the import to fail

For Apache Hive tables created with Apache Spark, there may be shell entities created in Apache Atlas. If these entities are in the export zip file before they are resolved, the import will fail.

Workaround: None

CDPD-70450: Impala SQL queries that include the “WITH” clause should populate lineage in Atlas

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-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-6565: Issue with ddlQueries in Atlas for Hive/Impala tables created in default/non-default db using DAS/HUE respectively

ddlQueries are not created for Impala origin table. ddlQueries created for Impala/Hive CTAS tables have difference in names as compared to same operation for Hive table. The Impala table has the query in the ddlQueries entity name.

None

CDPD-55301: ddlQueries and ALTERNATE_* lineage missing for Spark tables created through spark3-shell

The ddlQueries and ALTERNATIVE_* lineage missing for Spark tables created using spark3-shell.

None

CDPD-56085: 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

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*.

None

CDPD-58554: Discard audits of specific classification, label, and business metadata

Support to control audits for specific classification, label, and business metadata is not present in the Custom Audit Filters feature.

None

CDPD-40346: The ddlQueries and ALTERNATIVE_ADDCOLS lineage missing for Impala tables.

The ALTERNATIVE_ADDCOLS lineage has some issue when an Impala table is altered and the corresponding lineage is not created.

CDPD-67112: Import transforms do not work as expected when replacing a string which already has ":"

The character ":" is not supported in path replacements. The import succeeds but location remains unchanged. The character ":" must be avoided.

None

CDPD-69150: Unable to add labels or user defined properties in Japanese

Adding Japanese labels or user defined properties results in the error message: "Invalid label: ###, label should contain alphanumeric characters, _ or -"

None

CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties

Entities, such as hdfs_path, are not returned by the search when their indexable properties are searched using partial search terms made of Japanese or Chinese characters. Only exact matches return results when searching the indexed properties made of Japanese or Chinese characters.

None

CDPD-68191: Suggestions do not return the correct results when searching multiple Chinese characters

Free text search does not return results when searching Chinese phrases made of multiple characters. Partial searches return the correct results.

None.

CDPD-71219: Regression : Suggestions don't work for single character words on indexable attributes

When searching for entities whose name (entity names are indexable) is a single character, search results are returned but suggestions are not. When searching for entities whose description (entity descriptions are not indexable) is a single character, both search results and suggestions are returned.

None.

CDPD-67450: Table name renaming operation is not updating or creating iceberg_table entity

Renaming an Iceberg table does not update the corresponding Atlas entity.

None.

CDPD-67089: Export/Import: When a table with Ozone path is exported as "connected", only the Ozone key is exported.

When table with Ozone path is exported as "connected", only Ozone key is exported. Other Ozone entities, such as Ozone volume, Ozone bucket are not exported.

None.

CDPD-43772: Performance issues with Atlas service

If there are lot of update operations and the compression type of column families of atlas_janus table is SNAPPY, then the Kafka message processing might become slower.

- Consider setting compression type of column families of atlas_janus table as GZ.

OPSAPS-67783: During rolling upgrade one among two Atlas server failed to start but Cloudera Manager considered as success

Cloudera Manager marks the Execute command Start on service Atlas-1 as a success even when Atlas service had failed to start successfully. In such cases, Atlas logs give the exact reason for Atlas start-up failure.

None

CDPD-19358: "IsIndexable" and "isOptional" value of a typedef's attribute is modified post migration.

In HDP-265, falcon_feed_creation has "stored-in" attribute which has IsIndexable value set to True.

Post migration, "stored-in" is moved to relationshipAttributeDefs and has IsIndexable value set to False.

Similarly, in HDP-265 the table attribute of hive_storage_desc isOptional is set to true. Post migration, isOptional set to False.

None.

CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command

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

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

In converting content from Navigator to Atlas, the create time for Hive tables is not moved to Atlas.

None

CDPD-10574: Suggestion order doesn't match search weights

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

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@58112
bc4 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) does not 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 the metadata collection from HMS. For example, if an ALTER TABLE operation in Spark that is changing a table name and is reporting 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 "@" does not fetch the correct results

When searching Atlas qualifiedName values that include an "at" 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, and use 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 does 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 logout 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 CDP.

None.

CDPD-1664: Guest users are redirected incorrectly

Authenticated users logging in to Atlas are redirected to the CDP 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: The 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, to add metadata tags to relationships that you wanted to make sure were unique in the system, you can design the relationship attribute with the property IsUnique equal true. However, in this release, the IsUnique attribute is not enforced.

None.

DOCS-13759: Tag Propagation stops after a certain depth while the lineage is being extended

When a tag is added to an entity at timestamp T1, the entities along the lineage to which the tag must be propagated is calculated at T1. Before tag propagation completes, if the lineage is extended, tag does not propagate to the entities in the extended lineage.

CDPD-41142: When a Kafka console consumer group is run, more than one update audits are seen

After running the console consumer with a consumer group, verify the consumer group entity created, along with the metrics and notifications for the consumer group and topic. The expected result can be: one ENTITY_CREATE audit and one ENTITY_UPDATE audit. But more than one ENTITY_UPDATE audits are seen.

CDPD-40165: Two audits are created for SPARK CTAS table

When following Spark queries are fired:

```
spark.sql("create table table1(id int)")
spark.sql("create table table2 as select * from table1")
```

HMS sends "ENTITY_CREATE" and "ENTITY_FULL_UPDATE_V2".

The extra ENTITY_FULL_UPDATE_V2 message received from HMS is sent as a part of ALTER_TABLE_ADDCOLS event from the HMS Hook side. This behaviour is observed only when the queries are run from Spark SQL and not when run the same queries are run from Beeline.

CDPD-39197: Debug metrics returns empty data

When debug metrics is enabled and some operations are performed, the response is empty

CDPD-36495: Updating legacyAttribute from False to True resets the initially created relationshipAttributes values

Creating types, entities, and to start, you must set the relationship with is_legacy_attribute value as **False**.

Later, update the value relationshipDef is_legacy_attribute to True.

For the entities that were created before updating the is_legacy_attribute to True, the relationship Attributes value is reset.

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-29409: Hive import: Suggestion suggests entity which is deleted.

Suggestions suggests tables of a database, which is a deleted entity.

CDPD-25152: Tag propagation through deferred actions consumes additional time as compared to default flow

The additional time might be due to the small overhead added to create / update task vertex and which is run in the background. This also depends on number of tasks queued to be executed.

CDPD-42954: Zeppelin notebook fails after enabling Atlas-HDFS hook

The Zeppelin notebooks are failing with errors after enabling Atlas-HDFS hook in the CDP cluster.

When the below properties are set for atlas-client.properties in Cloudera Manager:

- atlas.jaas.KafkaClient.option.keyTab
- atlas.jaas.KafkaClient.option.principal

Along with adding the properties in /etc/atlas/conf/atlas-application.properties, Cloudera Manager also adds these properties to atlas-application.properties for other services (like Spark).

Adding these properties interferes with the normal flow of the services (like Spark)

To enable HDFS lineage feature, instead of setting these properties through Cloudera Manager, users can manually add the properties directly in /etc/atlas/conf/atlas-application.properties

CDPD-10576: Deleted Business Metadata attributes appear in Search Suggestions

Atlas search suggestions continue to show Business Metadata attributes even if the attributes have been deleted.

None.

CDPD-39427: [HDFS Lineage]: When the input is a directory in case of put/copyfromLocal/cp/mv commands, lineage is not created even though the script succeeds.

When Source is a directory and target is a directory which is already present in Atlas, the command succeeds and inserts the data in the desired location, but lineage is not created.

DOCS-13760: System Attributes search, __classificationNames: Search with parent tag does not return entities associated to its children tags

System attribute search with __classificationNames = parent_tag returns entities associated to parent_tag only and not entities associated to its children tag.

Workaround: Instead of using system attribute, employ the basic search attribute "classification" which lists entities associated with inherited classifications.

CDPD-35818: Basic search with tag filter provides approximateCount as -1 when there is no match and is 0 otherwise

When the following search operations are performed:

- Faceted search with both tag filter and entity filter

The observed approximateCount is -1.

Here both entity filter and tag filter are present and when there is no match the response received is -1.

- Faceted search with only entity filter

Performing a basic search provides an approximate value of 0 when there is no match.

- Faceted search with only tag filter

Whenever there is tag filter in the query and there is no entity match, the approximateCount is -1 and if the tag filter is not available, the response approximateCount is 0

CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

CDPD-22744: Bulk entity DELETE API does not delete authorised entities

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet `atlas.authentication.method.file` is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-82056: UI: when server response date fields as '0', UI shows as current time

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](#)

CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export

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

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](#)

CDPD-84502: Advanced Search is not working properly in Atlas UI

The Advance Search in Atlas does not return any results when selecting the Typedef option. The API is not triggered for each relation.

None

CDPD-76789: Creating tag with name description throws java.lang.ClassCastException

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

CDPD-70321: Atlas Parallel import is failing with various errors

During a parallel import-export activity with six iceberg table policies with exportOption as db1.* for all six exports, all import fail after the exports.

CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

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

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

CDPD-79099: Unable to import hive tables using import script in Public Cloud

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-88230: Bulk import operations currently do not construct lineage as expected

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

Bulk import operations do not create audits for the imported entities.

Known Issues in Cruise Control

Learn about the known issues in Cruise Control, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9

Rebalancing with Cruise Control does not work due to the metric reporter failing to report the CPU usage metric

On the Kafka broker, the Cruise control metric reporter plugin may fail to report the CPU usage metric.

If the CPU usage metric is not reported, the numValidWindows in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.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 the [Moving a Host Between Clusters](#).



Note: Cluster nodes affected by this issue are not displayed as unhealthy.

Known Issues in Hive

Learn about the known issues in Hive, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 9.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

CDPD-60770: Beeline Authentication Issue with Special Characters in Passwords

When LDAP is enabled, users cannot authenticate with Beeline if the password contains a special character. For example, the following string fails:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>;user=user@XXX;password='R3G#xpXyoy1M0Jb1'
```

Use the `-p` parameter to execute the Beeline command:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>; -n user@XXX -p 'R3G#xpXyoy1M0Jb1'
```

CDPD-43769: Inconsistent behavior with CHAR comparisons and string literals

In Hive, 'char' data to 'string' data comparison is considered as equal, but 'char' data to 'varchar' data comparison is considered as not equal if the data has trailing spaces. For example,

```
create table test(colchar char(10), colvarchar varchar(20), colstring string);
insert into test values ('a ', 'a ', 'a ');

select count(*) from test where colchar = colstring;
Output returns 1 because char and string data is considered as equal

select count(*) from test where colchar = colvarchar;
Output returns 0 because char and varchar data is considered as not equal
```

It is recommended that you use 'varchar or string' data type instead of 'char' data type for any character or string representation.

CDPD-42726: Wrong results for agg queries from stats after running multi-insert query on managed table

Multiple inserts at the same time into the same partition results in invalid stats. For example:

```
create table source(p int, key int, value string);
insert into source(p, key, value) values (101,42,'string42');

create table stats_part(key int,value string) partitioned by (p int);
from source
insert into stats_part select key, value, p
insert into stats_part select key, value, p;

select count(*) from stats_part;
-- In this scenario, StatsOptimizer helps serving this query because the result should be rowNum of the partition p=101.
-- The result is 1, however, the result should be 2.
```

Rewrite the multi-insert statement by merging the branches inserting into the same partition to one branch.

CDPD-57574: Query execution fails due to NullPointerException in DagUtils.setupQuickStart

Hive queries may fail at runtime with a `NullPointerException` while executing the Tez graph at `DagUtils.setupQuickStart` method. A part of the stack trace from the error is shown below:

```
ERROR : Failed to execute tez graph.
java.lang.NullPointerException: null at org.apache.hadoop.hive.ql.exec.tez.DagUtils.setupQuickStart(DagUtils.java:1724)
```

The error usually occurs from an invalid plan that has cycles created by semi-join or map-join edges.

Disable some of the optimizations in order to remove the cycles from the plan. It is sufficient to set one of the following properties:

- `set hive.tez.dynamic.partition.pruning=false;`
- `set hive.tez.dynamic.semijoin.reduction=false;`
- `set hive.auto.convert.join=false;`

To avoid any undesired performance regressions, disable the optimization(s) only for specific queries and not globally (through Cloudera Manager) for the whole workload or cluster.

CDPD-43107: SemanticException for INSERT INTO statement when Hive Cost-based Optimizer (CBO) is disabled

If you are running the `INSERT INTO` query and `hive.cbo.enable` is set to "false", the query fails with a `SemanticException`. For example,

```
set hive.cbo.enable=false;

CREATE TABLE mytable (
  id INT,
  str STRING
);

INSERT INTO mytable (id, str) VALUES (1, 'a');
```

```
Output:
org.apache.hadoop.hive.ql.parse.SemanticException: 0:0 Expected
2 columns for insclause-0/default@mytable; select produces 1 col
umns. Error encountered near token ''a''
```

This issue occurs only when columns are specified for the target table in the query. `INSERT INTO mytable values (1, 'a');` does not result in an exception.

Enable CBO and run the query.

CDPD-43957: HiveServer shuts down during replication due to high resource usage

During Hive replication (both bootstrap and incremental), you may notice that the `HiveServer` (HS2) shuts down periodically with the following error:

```
java.sql.SQLException: org.apache.hive.jdbc.ZooKeeperHiveClientException: Unable to read HiveServer2 configs from ZooKeeper
    at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.java:265)
    at org.apache.hive.jdbc.HiveDriver.connect(HiveDriver.java:107)
    at java.sql.DriverManager.getConnection(DriverManager.java:664)
    at java.sql.DriverManager.getConnection(DriverManager.java:247)
    at com.cloudera.enterprise.hive3qt.Hive3QueryTool$HiveOperation.execute(Hive3QueryTool.java:682)
    at com.cloudera.enterprise.hive3qt.Hive3QueryTool.main(Hive3QueryTool.java:935)
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Unable to read HiveServer2 configs from ZooKeeper
```

```

    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:177)
    at org.apache.hive.jdbc.Utills.configureConnParamsFromZooKeeper(Utills.java:580)
    at org.apache.hive.jdbc.Utills.parseURL(Utills.java:391)
    at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.java:263)
    ... 5 more
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Tried all existing HiveServer2 uris from ZooKeeper.
    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.getServerHosts(ZooKeeperHiveClientHelper.java:132)
    at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:172)
    ... 8 more

```

This issue occurs when you replicate at scale with an unbalanced cluster setup that has all the roles running on the Cloudera Manager host. As a result, Cloudera Manager ends up in a bottleneck situation because HS2 crashes, preventing further replication. The logs (/var/log/messages) indicate that there were a large number of Java processes running on the host, which exhausted the host's memory and triggered an Out Of Memory Killer process to stop HS2.

You can restart HS2 and reinitiate replication. However, the replication process may take a longer time to complete.

- Follow these recommendations when setting up your source and target clusters:
 - Ensure that no HS2, Hive metastore (HMS), DataNode, or NameNode roles are running on the Cloudera Manager host.
 - Ensure that you have multiple instances of HS2 and HMS roles on different nodes.
 - It is recommended that you have the HS2 and HMS roles on different nodes than the DataNodes or NameNodes.

These practices increase the possibility of Hive replication at scale completing successfully.

- Limit the scale of Hive replication by temporarily disabling replication policies.

CDPD-40730: Parquet change can cause incompatibility

Parquet files written by the parquet-mr library in CDP 7.1.9, where the schema contains a timestamp with no UTC conversion will not be compatible with older versions of Parquet readers. The effect is that the older versions will still consider these timestamps as they would require UTC conversions and will thus end up with a wrong result. You can encounter this problem only when you write Parquet-based tables using Hive, and tables have the non-default configuration `hive.parquet.write.int64.timestamp=true`.

None.

CDPD-41274: HWC + Oozie issue: Could not open client transport with JDBC Uri

Currently only Spark cluster mode is supported in the Oozie Spark Action with Hive Warehouse Connector (HWC).

Use Spark action in cluster mode.

```

<spark xmlns="uri:oozie:spark-action:1.0">
  ...
  <mode>cluster</mode>
  ...
</spark>

```

CDPD-45370: Cartesian join fails with an empty table

When the configuration `hive.tez.cartesian-product.enabled` is set to true, query fails if one of the tables involved in a Cartesian join has zero rows. This results in an `IllegalArgumentException` with the message `Number of items is 0. Should be positive.`

To address this issue, set `hive.tez.cartesian-product.enabled` to `false` before running the query.

Apache Jira: [\[HIVE-26582\]](#)

CDPD-26556 After an upgrade, querying a CTAS table under certain conditions might throw an exception

If you upgrade your Hive cluster from CDH 6 to CDP 7, create a CTAS table in the CDP cluster from a table you upgraded from CDH, you might see the following exception when you query the new table:

```
class org.apache.hadoop.io.IntWritable cannot be cast to class org.apache.hadoop.hive.serde2.objectinspector.StandardUnionObjectInspector$StandardUnion
```

This issue involves CDH-based tables having columns of complex types ARRAY, MAP, and STRUCT.

CDPD-23506: OutOfMemoryError in LLAP

Long running spark-shell applications can leave sessions in interactive Hiveserver2 until the Spark application finishes (user exists from spark-shell), causing memory pressure in case of a high number of queries in the same shell (1000+).

You must close spark-shell so that sessions are closed. Add the owner of the database or the tables as a user with read or read/write access to the tables directly.

CDPD-23041: DROP TABLE on a table having an index does not work

If you migrate a Hive table to CDP having an index, DROP TABLE does not drop the table. Hive no longer supports indexes ([HIVE-18448](#)). A foreign key constraint on the indexed table prevents dropping the table. Attempting to drop such a table results in the following error:

```
java.sql.BatchUpdateException: Cannot delete or update a parent row: a foreign key constraint fails ("hive"."IDXS", CONSTRAINT "IDXS_FK1" FOREIGN KEY ("ORIG_TBL_ID") REFERENCES "TBLS" ("TBL_ID"))
```

There are two workarounds:

- Drop the foreign key "IDXS_FK1" on the "IDXS" table within the metastore. You can also manually drop indexes, but do not cascade any drops because the IDXS table includes references to "TBLS".
- Launch an older version of Hive, such as Hive 2.3 that includes IDXS in the DDL, and then drop the indexes as described in [Language Manual Indexing](#).

Apache Issue: [Hive-24815](#)

CDPD-17766: Queries fail when using spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive.

Do not specify `spark.sql.hive.hiveserver2.jdbc.url.principal` in the JDBC URL to invoke Hive remotely.

Workaround: specify `principal=hive.server2.authentication.kerberos.principal` as shown in the following syntax:

```
jdbc:hive://<host>:<port>/<dbName>;principal=hive.server2.authentication.kerberos.principal;<otherSessionConfs>?<hiveConfs>#<hiveVars>
```

CDPD-13636: Hive job fails with OutOfMemory exception in the Azure DE cluster

Set the parameter `hive.optimize.sort.dynamic.partition.threshold=0`. Add this parameter in Cloudera Manager (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)

CDPD-10848: HiveServer Web UI displays incorrect data

If you enabled auto-TLS for TLS encryption, the HiveServer2 Web UI does not display the correct data in the following tables: Active Sessions, Open Queries, Last Max n Closed Queries

Known Issues in Hue

Learn about the known issues in Hue, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 9.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9**CDPD-88578: Missing Python Packages in Hue Resulting in Errors**

In Cloudera Runtime 7.1.9 SP1 CHF9, Hue encounters errors due to missing Python packages in its runtime environment. This leads to failures in Hue operations and thus you might experience failures or unexpected behaviour when executing Hue-related workflows or actions.

- Switch Hue to use Python 3.11 if not already configured. For more information, see [How to configure a specific Python version for Hue](#).
- Monitor Hue logs for missing package errors and manually install the required Python packages in the Hue virtual environment in the site-packages.

To install packages for Python 3.11, use the following command:

```
./build/env/bin/pip install <package-name>
```

For other Python versions (e.g., Python 3.8 or 3.9), use the following command:

```
./build/venvs/python3.x/bin/pip install <package-name>
```

- Upgrade to Cloudera Runtime 7.1.9 SP1 CHF10, where the issue is now resolved.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9**CDPD-58978: Batch query execution using Hue fails with Kerberos error**

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark queries in a batch mode launches a shell script Oozie job from Hue and this is not supported on a secure cluster.

There is no workaround. You can submit the queries individually.

OPSAPS-69659: Hue service fails on restart with "Unable to find pycopg2 2.5.4" error

Hue service fails to restart and you see the following error: Unable to find pycopg2 2.5.4. This could be because you have installed Python in a non-default location and Hue is unable to locate the pycopg2 PostgreSQL database adapter.

You must specify the path where you have installed Python in the PYTHONPATH property in the Hue Advanced Configuration Snippet using Cloudera Manager.

1. Log in to Cloudera Manager as an Administrator.
2. Go to Clusters Hue Configurations and add the following key and value in the Hue Service Environment Advanced Configuration Snippet (Safety Valve) field:

Key: PYTHONPATH

Value: `[***PYTHON-PATH***]`

Replace `[***PYTHON-PATH***]` with the actual location where you have installed Python. For example, `/opt/cloudera/parcels/CDH/lib/hue/build/env/lib/python3.8/site-packages`

3. Click Save Changes.
4. Restart the Hue service.

CDPD-59677: Unable to view Phoenix tables on the left assist in Hue

On clusters secured with Knox, you may not be able to see Phoenix tables on the left assist that are present under the default database (that is, an empty("") database).

None.

CDPD-58142: A query is not pre-populated in the Hue editor after clicking on the “Re Execute” button

When you click Re Execute to rerun a query from the Job Browser Queries Query Details page, the query does not get populated on the Hue editor, as expected.

None.

CDPD-54376: Clicking the home button on the File Browser page redirects to HDFS user directory

When you are previewing a file on any supported filesystem, such as S3 or ABFS, and you click on the Home button, you are redirected to the HDFS user home directory instead of the user home directory on the said filesystem.

None.

CDPD-41306: pip3.8 freeze command does not work and results into an error

You see the `/usr/lib/hue/build/env/bin/python: No such file or directory` error when you run the following command:

```
build/env/bin/pip3.8 freeze
```

Run the freeze command as follows by specifying the paths of python3.8 and pip3.8:

```
/opt/cloudera/parcels/CDH/lib/hue/build/env/bin/python3.8 /opt/cloudera/parcels/CDH/lib/hue/build/env/bin/pip3.8 freeze
```

CDPD-43293: Unable to import Impala table using Importer

Creating Impala tables using the Hue Importer may fail.

If you have both Hive and Impala services installed on your cluster, then you can import the table using by selecting the Hive dialect from Tables Sources .

If only Impala service is installed on your cluster, then go to Cloudera Manager Clusters Hue Configurations and add the following line in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini field:

```
[beeswax]
max_number_of_sessions=1
```

DWX-8602: Unable to import a large CSV file from the local workstation

You may see an error message while importing a CSV file into Hue from your workstation, stating that you cannot import files of size more than 200 KB.

Upload the file to S3 or ABFS and then import them into Hue using the Importer.

CDPD-39330: Unable to use the pip command in CDP

You may not be able to use the pip command in CDP 7.1.7 or higher and may see the following error when using pip in a command: “ImportError: cannot import name chardet”.

Follow the steps listed on [Unable to use pip command in CDP](#).

CDPD-24294: Hue uses the unsafe-inline directive in its Content Security Policy (CSP) header

Hue 4 web interface uses the unsafe-inline directive in its CSP header. As a result, the application server does not set the CSP header in its HTTP responses, and therefore does not benefit from the additional protection against potential cross-site scripting issues and other modern application vulnerabilities which a properly configured CSP may provide. This could lead to application vulnerability.

Cloudera recommends deploying additional security measures such as a firewall within the Hue server to control allowed connections, and SSO-based authentications mechanisms such as LDAP or SAML.

OPSAPS-58927: Connection failed error when accessing the Search app (Solr) from Hue

If you are using Solr with Hue to generate interactive dashboards and for indexing data, and if you have deployed two Solr services on your cluster and selected the second one as a dependency for Hue, then Cloudera Manager assigns the hostname of the first Solr service and the port number of the second Solr service generating an incorrect Solr URL in the search section of the hue.ini file. As a result, you may see a “Connection failed” error when you try to access the Search app from the Hue web UI.

1. Log into Cloudera Manager as an Administrator.
2. Go to Clusters Hue service Configuration and add the following lines in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini field:

```
[search]
# URL of the Solr Server
solr_url=http://[***HOSTNAME***]:[***PORT***]/solr/
```

For example:

```
solr_url=http://solr2:4567/solr/
```

3. Click Save Changes.
4. Restart the Hue service.

CDPD-41136: Importing files from the local workstation is disabled by default

Cloudera has disabled the functionality to import files from your local workstation into Hue because it may cause errors. You may not see the Local File option in the Type drop-down menu on the **Importer** page by default.

You can enable the functionality to import files from your local workstation by specifying the following parameter in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue_safety_valve.ini field in Hue configurations in Cloudera Manager:

```
[indexer]
enable_direct_upload=true
```

Invalid S3 URI error while accessing S3 bucket

The Hue Load Balancer merges the double slashes (//) in the S3 URI into a single slash (/) so that the URI prefix "/filebrowser/view=S3A://" is changed to "/filebrowser/view=S3A:". This results in an error when you try to access the S3 buckets from the Hue File Browser through the port 8889.

The Hue web UI displays the following error: “Unknown error occurred”.

The Hue server logs record the “ValueError: Invalid S3 URI: S3A” error.

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.

You should be able to load the S3 browser from both 8888 and 8889 ports.

Alternatively, you can use the Hue server port 8888 instead of the load balancer port 8889 to resolve this issue.

Error while rerunning Oozie workflow

You may see an error such as the following while rerunning an already executed and finished Oozie workflow through the Hue web interface: E0504: App directory [hdfs://cdh/user/hue/oozie/workspaces/hue-oozie-1571929263.84] does not exist.

To resolve this issue, add the following property in the Hue Load Balancer Advanced Configuration Snippet:

1. Sign in to Cloudera Manager as an administrator.
2. Go to Clusters Hue service Configurations Load Balancer and search for the Load Balancer Advanced Configuration Snippet (Safety Valve) for httpd.conf field.
3. Specify MergeSlashes OFF in the Load Balancer Advanced Configuration Snippet (Safety Valve) for httpd.conf field.
4. Click Save Changes.
5. Restart the Hue Load Balancer.

CDPD-17956, CDPD-14167: Setting idle session timeout for Hue does not work when the cluster is secured using Knox SSO

If Hue is configured with desktop.auth.backend.KnoxSpnegoDjangoBackend as the Authentication Backend, then the automatic idle session logout that is set by configuring the `idle_session_timeout` property does not take effect. You may also see 404 error while accessing Hue from the Knox UI when the `idle_session_timeout` property is not set to -1.

None

Query history displays "Result Expired" message

You see the "Result Expired" message under the Query History column on the **Queries** tab for queries which were run back to back. This is a known behaviour.

None.

CDPD-43956: Manually replace UUID when importing Oozie workflows containing sub-workflows

When importing Oozie workflows that contain sub-workflows, you must replace all the UUID entries with unique new entries. If you change the UUID of a sub-workflow, you must update that reference in the parent workflow to avoid circular dependencies.

Hue limitation after upgrading from CDH to CDP Private Cloud Base

The `hive.server2.parallel.ops.in.session` configuration property changes from TRUE to FALSE after upgrading from CDH to CDP Private Cloud Base. Current versions of Hue are compatible with this property change; however, if you still would like to use an earlier version of Hue that was not compatible with this property being FALSE and shared a single JDBC connection to issue queries concurrently, the connection will no longer work after upgrading.

Hue UI is blank upon login after upgrading to CDP 7.1.7 from CDH 6

If your cluster was secured using Knox, and if you have upgraded from CDH 6 to CDP 7.1.7, then you may see a blank Hue screen. This could happen because the `knox_proxyhosts` parameter is newly introduced in CDP, and it is possible that this parameter is not configured in Cloudera Manager under Hue configuration.

Specify the host on which you have installed Knox in the Hue Knox Proxy Hosts configuration as follows:

1. Log in to Cloudera Manager as an Administrator.
2. Obtain the host name of the Knox Gateway by going to Clusters Knox service Instances .

3. Go to Clusters Hue service Configuration and search for the Knox Proxy Hosts field.



Note: Cloudera Manager displays the following warning if the Knox Proxy Hosts field is empty when Knox Gateway is enabled on the CDP cluster: The parameter `knox_proxyhosts` cannot be empty. This can happen if there are no Knox Gateways. Please set the `knox_proxyhosts` to the list of hosts that have Knox Gateways.

4. Specify the Knox Gateway hostname in the Knox Proxy Hosts field.
5. Click Save Changes and restart the Hue service.

Unsupported features

CDPD-59595: Spark SQL does not work with all Livy servers that are configured for High Availability

SparkSQL support in Hue with Livy servers in HA mode is not supported. Hue does not automatically connect to one of the Livy servers. You must specify the Livy server in the Hue Advanced Configuration Snippet as follows:

```
[desktop]
[spark]
livy_server_url=http(s)://[***LIVY-FOR-SPARK3-SERVER-HOST***]:
[***LIVY-FOR-SPARK3-SERVER-PORT***]
```

Moreover, you may see the following error in Hue when you submit a SparkSQL query: Expecting value: line 2 column 1 (char 1). This happens when the Livy server does not respond to the request from Hue.

Specify all different Livy servers in the `livy_server_url` property one at a time and use the one which does not cause the issue.

CDPD-18491: PySpark and SparkSQL are not supported with Livy in Hue

Hue does not support configuring and using PySpark and SparkSQL with Livy in CDP Private Cloud Base.

Importing and exporting Oozie workflows across clusters and between different CDH versions is not supported

You can export Oozie workflows, schedules, and bundles from Hue and import them only within the same cluster if the cluster is unchanged. You can migrate bundle and coordinator jobs with their workflows only if their arguments have not changed between the old and the new cluster. For example, hostnames, NameNode, Resource Manager names, YARN queue names, and all the other parameters defined in the `workflow.xml` and `job.properties` files.

Using the import-export feature to migrate data between clusters is not recommended. To migrate data between different versions of CDH, for example, from CDH 5 to CDP 7, you must take the dump of the Hue database on the old cluster, restore it on the new cluster, and set up the database in the new environment. Also, the authentication method on the old and the new cluster should be the same because the Oozie workflows are tied to a user ID, and the exact user ID needs to be present in the new environment so that when a user logs into Hue, they can access their respective workflows.



Note: Migrating Oozie workflows from HDP clusters is not supported.

Known Issues in Kafka

Learn about the known issues in Kafka, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9**CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress**

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker) in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Save Changes > Restart SMM.

The offsets.topic.replication.factor property must be less than or equal to the number of live brokers

The offsets.topic.replication.factor broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a GROUP_COORDINATOR_NOT_AVAILABLE error until the cluster size meets this replication factor requirement.

None

Requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true

The first few produce requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true.

Increase the number of retries in the producer configuration setting retries.

Performance degradation when SSL Is enabled

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with ssl.secure.random.implementation = SHA1PRNG. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

OPSAPS-43236: Kafka garbage collection logs are written to the process directory

By default Kafka garbage collection logs are written to the agent process directory. Changing the default path for these log files is currently unsupported.

None

RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transactional method because we are in an error state
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeFailWithError(TransactionManager.java:1125)
    at org.apache.kafka.clients.producer.internals.TransactionManager.maybeAddPartition(TransactionManager.java:442)
    at org.apache.kafka.clients.producer.KafkaProducer.doSend(KafkaProducer.java:1000)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:914)
    at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:800)
    .
    .
    .
    Caused by: org.apache.kafka.common.errors.ClusterAuthorizationException: Cluster authorization failed.
```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting `enable.idempotence` to `false`.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

DBZ-4990: The Debezium Db2 Source connector does not support schema evolution

The Debezium Db2 Source connector does not support the evolution (updates) of schemas. In addition, schema change events are not emitted to the schema change topic if there is a change in the schema of a table that is in capture mode. For more information, see [DBZ-4990](#).

None.

CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy compression. The HDFS Stateless Sink connector is only affected if the `Compression Codec` or `Compression Codec` for Parquet properties are set to `SNAPPY`.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: org.apache.hadoop.util.NativeCodeLoader.buildSupportsSnappy()
```

```
Failed to write to HDFS due to java.lang.RuntimeException: native snappy library not available: this version of libhadoop was built without snappy support.
```

Download and deploy missing libraries.



Important: Ensure that you complete steps 1-11 on all Kafka Connect hosts. Additionally, ensure that the advanced configuration snippet in step 12 is configured for all Kafka Connect role instances.

1. Create the /opt/nativelibs directory.

```
mkdir /opt/nativelibs
```

2. Change the owner to kafka.

```
chown kafka:kafka /opt/nativelibs
```

3. Locate the directory containing the Hadoop native libraries and copy its contents to the directory you created.

```
cp /opt/cloudera/parcels/CDH/lib/hadoop/lib/native/* /opt/nativelibs
```

4. Verify that 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

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 Knox

Learn about the known issues in Knox, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 9.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

There are no known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9

There are no known issues identified in earlier releases.

Known Issues in Kudu

Learn about the known issues in Kudu, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 9.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9

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.

Workaround: Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

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.

Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

Unsupported features

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 in Cloudera Runtime 7.1.9 SP1 CHF 9.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

The following known issues were identified in this release:

No new known issues in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9

No known issues identified in Livy.

Known Issues in Iceberg

Learn about the known issues in Iceberg, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 9.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9

There are no known issues identified in earlier releases.

Known Issues in Impala

Learn about the known issues in Impala, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 9.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

CDPD-87697 : Python 3.11 Support in impala-shell

Python 3.11 is not supported in the CDP 7.1.9 SP1 CHF9 versions of the `impala-shell`

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the `--tls_ciphersuites` startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.

Ensure that the list in the `--tls_ciphersuites` startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

Apache Jira: [IMPALA-14625](#)

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work

If you try to run `impala-shell` on RHEL 9 by setting the default python executable available in `PATH` to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, `IMPALA_PYTHON_EXECUTABLE=python3`.

CDPD-81980: Impala query may fail due to backpressure

Impala queries could fail with a dropped due to backpressure error, even when service queue and memory usage appear normal.

Increase the value of the `--mem_limit` startup option.

CDPD-42958: After upgrading the CDH 7.1.9 from CDH 6.x, under certain conditions you cannot insert data into a table

Under the following conditions, after upgrading from CDH 6.x to CDH 7.1.9 you cannot insert data into a table from Impala:

- On CDH 6.x, you created a database with Impala in a user specified HDFS location.
- Using Hive, you then created a table in the database.

Under these conditions, the database and table are stored in the user-specified HDFS directory. After upgrading, the HDFS directory of the table is read-only for Impala. Consequently, from Impala you cannot insert new data into the table because Impala does not have write permission on the HDFS directory.

Workaround: To resolve this issue, use either one of the following workarounds:

- Using the Ranger Web UI, in the policy repository `cm_hdfs`, grant the user 'impala' write permission on the directory where the table resides.
- Enter the following command to grant write permission to user 'impala' on the HDFS directory where the table resides.

```
hdfs dfs -setfacl -m default:user:impala:rwX <HDFS directory>
```

CDPD-80019: Impala fails to start due to large catalog-update topic

Impala coordinators fail to start if the catalog-update statestore topic exceeds 2GB, hitting the Thrift 0.16.0 message size limit (`thrift_rpc_max_message_size`). This prevents the coordinator from processing the update.

- Enable local catalog mode: `--use_local_catalog=true`
- In Cloudera Manager, go to Impala Configuration search for "Catalog Server Command Line

Argument Advanced Configuration Snippet (Safety Valve)" and click 

- Add property name and value:

```
Name : --catalog_topic_mode
Value: minimal
```

- Save changes.

Apache Jira: [IMPALA-12350](#)

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 HDF
S file hdfs://nameservice1/warehouse/tablespace/managed/hive/<da
tabase>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException:
File does not exist: /warehouse/tablespace/managed/hive/<data
base>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

You must use a single Knox node to access Impala UI.

Impala api calls via 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-21828: Multiple permission assignment through grant is not working

None

Problem configuring masking on tables using Ranger

The following Knowledge Base article describes the behavior when we configure masking on tables using Ranger. This configuration works for Hive, but breaks queries in some scenarios for Impala.

For a workaround, see the following Knowledge Base article: [ERROR: "AnalysisException: No matching function with signature: mask\(FLOAT\)" when Impala jobs fail with the following error with signature: mask\(FLOAT\)](#)

IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXXXX"

IMPALA-532: Impala should tolerate bad locale settings

If the LC_* environment variables specify an unsupported locale, Impala does not start.

Add LC_ALL="C" to the environment settings for both the Impala daemon and the Statestore daemon.

Avro Scanner fails to parse some schemas

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

IMPALA-691: Process mem limit does not account for the JVM's memory usage

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the impalad daemon.

To monitor overall memory usage, use the `top` command, or add the memory figures in the Impala web UI `/memz` tab to JVM memory usage shown on the `/metrics` tab.

IMPALA-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Remove trailing semicolon from the Avro schema.

IMPALA-1652: Incorrect results with basic predicate on CHAR typed column

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the `RPAD()` function to blank-pad literals compared with `CHAR` columns to the expected length.

IMPALA-1821: Casting scenarios with invalid/inconsistent results

Using a `CAST()` function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

IMPALA-2005: A failed CTAS does not drop the table if the insert fails

If a `CREATE TABLE AS SELECT` operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

Drop the new table manually after a failed CREATE TABLE AS SELECT

IMPALA-2422: % escaping does not work correctly when occurs at the end in a LIKE clause

If the final character in the RHS argument of a LIKE operator is an escaped \% character, it does not match a % final character of the LHS argument.

None

IMPALA-2603: Crash: impala::Coordinator::ValidateCollectionSlots

A query could encounter a serious error if includes multiple nested levels of INNER JOIN clauses involving subqueries.

None

IMPALA-3094: Incorrect result due to constant evaluation in query with outer join

An OUTER JOIN query could omit some expected result rows due to a constant such as FALSE in another join clause. For example:

```
explain SELECT 1 FROM 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;
+---+
| id | select_type | table | partitions | type |
+---+
| 1  | SIMPLE     | a1    |              | eq_j |
| 2  | SIMPLE     | a2    |              | eq_j |
| 3  | SIMPLE     | a3    |              | eq_j |
+---+
```

```
| Explain String |  
+-----+  
| Estimated Per-Host Requirements: Memory=1.00KB VCores=1 |  
| 00:EMPTYSET |  
+-----+
```

IMPALA-3509: Breakpad minidumps can be very large when the thread count is high

The size of the breakpoint minidump files grows linearly with the number of threads. By default, each thread adds 8 KB to the minidump size. Minidump files could consume significant disk space when the daemons have a high number of threads.

Add `-minidump_size_limit_hint_kb=size` to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then 2 KB per thread after that.) The minidump file can still grow larger than the "hinted" size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

IMPALA-4978: Impala requires FQDN from hostname command on Kerberized clusters

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the `gethostname()` system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a Kerberized cluster.

Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname`, only returns the short name, pass the command-line flag `##hostname=FULLY_QUALIFIED_DOMAIN_NAME` in the startup options of all Impala-related daemons.

IMPALA-6671: Metadata operations block read-only operations on unrelated tables

Metadata operations that change the state of a table, like `COMPUTE STATS` or `ALTER RECOVER PARTITIONS`, may delay metadata propagation of unrelated unloaded tables triggered by statements like `DESCRIBE` or `SELECT` queries.

None

IMPALA-7072: Impala does not support Heimdal Kerberos

None

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.

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.

Known issues in Navigator Encrypt

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround.

Failed to start navencrypt-mount.service

After installing NavEncrypt on SLES, if the command "systemctl status navencrypt-mount" fails with the error: "Failed to start navencrypt-mount.service: Unit navencrypt-mount.service failed to load: No such file or directory", the systemd files need to be reloaded.

Run the command:

```
(sudo) systemctl daemon-reload
```

Crash in Navigator Encrypt kernel module if --profile-file= option is used with ACLs

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 Ozone

Learn about the known issues in Ozone, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 9.

Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9

CDPD-84518: Expired secret key can abort leader Ozone Manager startup

Ozone Manager fails to startup due to NPE exception because of expired SCM managed secret key.

Apache Jira: [HDDS-13234](#)

CDPD-64865: Intermittent timeout in TestBlockDeletion.testBlockDeletion (BlockDeletionService stuck)

If the event of SCM exiting safe mode was triggered multiple times, this might eventually cause the SCMBlockDeletingService to transition into the PAUSING state.

Check the SCM logs and ensure that the SCMBlockDeletingService is not running, and then restart the SCM.

CDPD-81569: Recon is not reflecting correct metrics for Volumes, Buckets, Keys, and Pending Key Deletions

Recon is displaying data for container and space utilisation only. The metrics for Volumes, Buckets, Keys, and Pending Key Deletions are not being updated and displayed correctly.

Apache Jira: [HDDS-12607](#)

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: 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.

The Prometheus binary is not available in CDP Private Cloud Base 7.1.9 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.

After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and if Ozone is in the Non-HA environment, an exception message is observed during the finalization of the Ozone upgrade.

During the finalization of the upgrade, a `ClassNotFoundException` message for `org.cloudera.log4j.redactor.RedactorAppender` class may be displayed. The error message can be ignored because the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

CDPD-68951: In CDP Private Cloud Base 7.1.9 CHF2 version and lower, the command `ozone sh key list <bucket_path>` displays the `isFile` flag in a key's metadata as `false` even when the key is a file. This issue is rectified in CDP Private Cloud Base 7.1.9 CHF3. However, the pre-existing (pre-upgrade) key's metadata cannot be changed.

None

CDPD-67771: When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(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.

CDPD-64398: In SCM with non-HA configuration, Secret key manager does not initialise. Hence, the startup of OM and DataNode fails because it cannot get a secret key. This key is used when security is enabled and used in Block token and container token verification while communicating between the Ozone client, OM, and Datanode.

This issue does not occur in SCM with HA configuration.

You must force exit from the safe mode for SCM. This triggers the initialization of the Secret key manager. Below are the options:

- exit safe mode manually whenever SCM is started
- disable safe mode by setting `hdds.scm.safemode.enabled=false` in safety valve for SCM configuration

Impact of disabling safe mode and the purpose of safe mode:

- Write should not fail once the cluster is out of safe mode
- Read of existing data should not fail after the cluster is out of safe mode
- Unnecessary re-replication should be avoided during cluster restart

So disabling safe mode does not have any major impact on Ozone.

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.

CDPD-52412: keyManagerImpl#listStatus exceeds the maximum RPC length.

To reduce the payload, you must increase the part size.

OPSAPS-68159: If you did not deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Error when distributing parcel to host, Ignoring non-compliant parcel manifest error is displayed after Cloudera Manager is upgraded to 7.11.3.

If you encounter the error, perform the following steps:

1. Deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.11.3.
2. Restart the cluster with a delay of 10 minutes.
3. Continue to upgrade CDH 7.1.8 to CDH 7.1.9.

OPSAPS-68159: If you did not deactivate the Ozone parcel 718.2.x on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Ozone roles during the CDH 7.1.8 upgrade to CDH 7.1.9.

If you encounter the error, perform the following steps:

1. Deactivate the Ozone parcel 718.2.x.
2. Restart the Ozone service.
3. Perform Finalize Upgrade for the Ozone service.

Step result: The Ozone roles will be displayed in green.

CDPD-60989: The packaging version for Apache Ozone points to the 1.2.0 older version. This is a version string problem and not a packaging issue. The version of the Apache Ozone binary is closest to 1.4.0.

None. This only affects the JAR versioning.

CDPD-60366: Native library loader fails when system property `native.lib.tmp.dir` is not set. It fails because the library is copied to `/` instead of the `cwd`.

Setting `native.lib.tmp.dir` to a certain path like `/tmp` should solve the issue. To set this system property, you must add `-Dnative.lib.tmp.dir=/tmp` to `ozone_java_opts` in Cloudera Manager configuration of the Ozone cluster.

CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

CDPD-60598: DataNode can consume up to the JVM heap size worth of direct memory buffers under specific failure scenarios. This can lead to a DataNode crash. This is due to Netty allocating and not freeing the direct memory buffers.

Restarting DataNode will resolve this issue.

CDPD-60466: The Ozone client is missing in the Cloudera Manager's Service Monitor node. This causes Ozone Canary's health check to fail.

Ensure you install Ozone CLI on the same node as Cloudera Manager's Service Monitor.

CDPD-60012: The Ozone SCM may be stuck in safe mode after upgrade to 7.1.9. This is due to a bug in how the SCM accounts for open storage containers.

The workaround is to temporarily set the Ozone configuration `hdds.scm.safemode.threshold.pct` to a lower value like 0.90 and restart the SCM.

CDPD-59679: In some scenarios, excess UNHEALTHY replicas of EC containers may not be removed from the cluster.

None

CDPD-50116: Topology-aware reads can provide better performance by routing applications to the closest replica of a file, if available, on the same physical node or same rack. However this feature is disabled by default in CDP 7.1.9.

This feature can be enabled by setting `ozone.network.topology.aware.read` to true using Cloudera Manager safety valve.

CDPD-57165: The DataNode disk usage thread may abort due to Ratis log tailing.

Use `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory` to calculate disk usage. Configure in the `ozone-site.xml` property: `hdds.datanode.du.factory` value: `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory`

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.

CDPD-55513: Hive external table replication policy does not migrate Hive tables in HDFS storage to Ozone.

To replicate data from HDFS to Ozone, see [Migrating your data from HDFS to Ozone](#). You can use HMS Mirror to replicate the metadata.

OPSAPS-67757: Hive external tables in Ozone storage cannot be replicated using Hive external table replication policies.

To replicate the Hive external tables' data, consider using DistCp. To replicate the metadata of Hive external tables, consider using HMS Mirror.

Remove bucket recursively using `rb --force` command from AWS S3 cannot work for FSO buckets.

Use the Ozone shell command `ozone sh bucket delete -r [***BUCKET ADDRESS***]`

CDPD-59126: Info log displays "noexec permission on /tmp/liborg_apache_ratis_thirdparty_netty_transport_native_epoll_x86" on client while executing command with noexec on /tmp.

To suppress Info log related to `liborg_apache_ratis_thirdparty_netty_transport_native_epoll_x86` library: Export `OZONE_OPTS` environment variable on the client terminal by running the command `export OZONE_OPTS="-Dorg.apache.ratis.thirdparty.io.netty.native.workdir=/var/tmp $OZONE_OPTS"`

OPSAPS-67650: Ozone uses RocksDB as a library to persist metadata locally.

By default, RocksDB places certain executables in `/tmp`, and thus encounters errors when `/tmp` is mounted with `noexec`.

The workaround is to configure RocksDB to put executables at another location. On a PhatCat node, the steps are:

1. Go to Cloudera Manager UI OZONE Configuration .
2. Find Ozone Service Environment Advanced Configuration Snippet (Safety Valve) and set the following environment variable: `ROCKSDB_SHAREDLIB_DIR=/VAR/TMP`
3. Restart Ozone.

CDPD-49137: Ozone Manager Kerberos token expires for SCM communication and OM does not log in again.

Sometimes, OM's Kerberos token is not updated and it stops to communicate with SCM. When this occurs, writes start failing.

Restart OM or set the safety valve `hadoop.kerberos.keytab.login.autorenewal.enabled = true`

CDPD-56684: Keys get deleted when you do not have permission on volume

When a volume is deleted, it recursively deletes the buckets and keys inside it and only then deletes the volume. The volume delete ACL check is done only in the end, due to which you may end up deleting all the data inside the volume without having delete permission on the volume.



Note: There was no bucket/key permission available which allowed the user to delete them recursively.

CDPD-50610: Large file uploads are slow with OPEN and stream data approach

Hue file browser uses the append operation for large files. This API is not supported by Ozone in 7.1.9 and therefore large file uploads can be slow or timeout from the browser.

Use native Ozone client to upload large files instead of the Hue file browser.

OPSAPS-64097: Ozone service restart failed at SCM

Stopping SCM service using Cloudera Manager can sometimes timeout and need a retry. The Cloudera Manager API waits for 90 seconds which is not sufficient under certain circumstances.

Retry the shutdown using Cloudera Manager if the SCM still shows as running after a refresh of the service status.

OPSAPS-66469: Ozone-site.xml is missing if the host does not contain HDFS roles

The client side `ozone-site.xml` (`/etc/hadoop/conf/ozone-site.xml`) is not generated by Cloudera Manager if the host does not have any HDFS role. Because of this, issuing Ozone commands from

that host fails because it cannot find the service name to host name mapping. The error message is similar to this: # ozone sh volume list o3://ozoneabc 23/03/06 18:46:15 WARN ha.OMProxyInfo: OzoneManager address ozoneabc:9862 for serviceID null remains unresolved for node ID null Check your ozone-site.xml file to ensure ozone manager addresses are configured properly.

Add the HDFS gateway role on that host.

OPSAPS-67607: Cloudera Manager FirstRun failure at the “Upload YARN MapReduce Framework JARs” step.

If this failure is attributed to the broken symbolic link, /var/lib/hadoop-hdfs/ozone-file-system-hadoop3.jar, it is likely due to the presence of the user hdfs on the node prior to CDP parcel activation. As a result, the Cloudera Manager agent skips the initialization related to HDFS, leading to the non-creation of the /var/lib/hadoop-hdfs directory.

Create the directory “/var/lib/hadoop-hdfs” on all nodes followed by the deactivation and activation of the CDP parcel (deactivate and activate the Ozone parcel instead, in case Ozone parcel is used).

CDPD-50447: When SCM High Availability is enabled, each of the SCM web UIs report the host of the web UI as the leader of HA, and the other two as followers. This gives wrong information

Correct output is available by running the `ozone admin scm roles --service-id=<ID>` command.

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](#).

OPSAPS-62327: In an Ozone cluster without any gateway roles, Ozone is unable to deploy client configurations and displays the ConfigGenException error.

You must add the Ozone gateway roles to the cluster.

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-36389: The configurations `datanodes.involved.max.percentage.per.iteration` and `size.moved.max.per.iteration` are meant to limit the max number of DataNodes that is involved and maximum size that can move in an iteration. This bug causes balancer to stop an iteration when it is two DataNodes or 1 Container size (5GB) away from hitting these limits. However, these DataNodes can again be considered for balancing in the next iteration. This means the cluster is balanced after enough iterations, albeit a bit slowly. This bug is apparent in small clusters of around four DataNodes where the DataNodes can be either the source or target for a lot of moves but the iteration gets stopped when three DataNodes have been involved. It takes a higher number of iterations to eventually balance this cluster. While this is a performance issue, it does not prevent balancer from ultimately balancing the cluster.**

Increase the speed for balancing by decreasing the interval between each iteration using the configuration `balancing.iteration.interval`. The value of this configuration must be greater than `hdds.datanode.du.refresh.period`. `size.moved.max.per.iteration` can be increased to allow more data to move in one iteration.

**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-35141: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2> (state=08S01,code=40000) java.sql.SQLException: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2>. This issue does not occur if the source and target buckets are different in Hive queries. Currently, copying across the same bucket is only supported.**

Avoid different buckets in source and target path. Ensure that `hive.exec.stagingdir` is unset or set to the same bucket as the warehouse so that staging directories are created in the same bucket as the warehouse.

**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-41184: With Legacy buckets, FileSystem Optimised is not interoperating with the Ozone shell command. The cause is that the directory key entry in the DB KeyTable is stored as "dir1/" with trailing slash. But while performing the described operation, Ozone shell (o3://) is normalizing the given path and removing the trailing slash "/" from it. That results in KEY\_NOT\_FOUND exception.**

There are three workarounds:

- Use FileSystem API to Delete the Directories rather than Shell-Command API.
- Use FSO buckets instead of Legacy Buckets. As in FSO, you can create Intermediate Directories and Delete Directories using the Ozone shell commands.
- Disable and set the configuration `ozone.om.enable.filesystem.pathsflag` to false in order to delete the directories. This is generally not a preferred workaround because the cluster must be restarted to pick up the new changes.

**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-12542: Mount of Ozone filesystem with the help of FUSE fails.**

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-42897: EC writes are failing with "No enough datanodes to choose" after EC replication config is set globally.**

EC writes start failing when large number of pipelines are created as a result of multiple EC configs with different chunk sizes used to write keys.

If standard EC configs (that is., rs-3-2-1024k) are used to write keys, number of pipelines created per DataNode is limited to five and this issue is not seen with standard EC configs.

The recommendation is not to create too many random chunk sizes. It is configurable because, users can decide based on their workload. But do not to have separate chunk sizes for each file.

**CDPD-41539: The "No such file or directory" imessage is displayed when EC file is read using older of client.**

You must upgrade the client before trying to read the key: vol1/ecbuck1/1GB\_ec".

**CDPD-40560: Filesystem Operations using Hadoop S3a connector on a FILE\_SYSTEM\_OPTIMIZED bucket fails. org.apache.hadoop.ozone.om.exceptions.OMException: Unable to get file status: volume: s3v bucket: fso key: test/**

Do not run Hadoop S3a commands on a FILE\_SYSTEM\_OPTIMIZED bucket. Use OBJECT\_STORE bucket layouts.

**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.

**CDPD-43432: Ozone Service in fault state in DataNode - Long Running setup.**

Upgrade RocksDB to the latest version.

**OPSAPS-63999: In the newly installed cluster, the Finish upgrade option is clickable.**

None

**OPSAPS-64648: Failed to start Ozone node using Cloudera Manager if default log path /var/log/hadoop-ozone does not exist. If this path does not exist, any Ozone nodes(for example SCM or DataNode) restart will fail.**

Run the following command `sudo -u hdfs mkdir -p /var/log/hadoop-ozone` or replace hdfs with the user Ozone roles that are running.

**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.

### Known Issues in Ranger

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 9.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

There are no new known issues identified in this release.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

#### UnsupportedClassVersionError

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.

#### CDPD-68806: The Revoke operation for users belonging to a group or role permission does not function as expected.

List command is listing all the tables even when the user permission is revoked. And also the command does not add any deny policy to Ranger for that specific user.

For example:

user\_1 is the part of group\_1

user\_2 is the part of role\_1

Create a Ranger policy for group\_1 and role\_1 to access all the resources in HBase.

Try to revoke the permission for user\_1 and user\_2 using HBase shell global revoke command:

```
revoke 'user_1'
```

```
revoke 'user_2'
```

But there is no deny policy added for that user. And access is still granted to user\_1 and user\_2.

This behavior is currently not supported in HBase shell. Must be handled manually using the Ranger policy change.

#### CDPD-68739: The revoke command does not work when using the HBase shell

While using the HBase shell, running the revoke command does not cancel the user permission. Users are able to perform actions even after running the revoke command.

For example:

When a user is granted permission to a particular table using the HBase grant command.

```
hbase:003:0> grant 'user_1', 'RWCA', 'some_abc_table'
```

And then execute the global revoke command for that user:

```
hbase:005:0> revoke 'user_1'
```

Permission for that user 'user\_1' for the table 'some\_abc\_table' not got revoked.

None.

#### CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

#### CDPD-58860: After upgrading CDP-7.1.8 to CDP-7.1.9, cdp-proxy token missing from Knox-Ranger policy.

As part of [OPSAPS-67480](#) in 719 default ranger policy is added from cdp-proxy-token topology, so that after a new installation of CDP-7.1.9, the Knox-ranger policy includes cdp-proxy-token. However, upgrades do not add cdp-proxy-token to cm\_knox policies automatically.

Manually add cdp-proxy-token to the Knox policy, using Ranger Admin Web UI.

1. Log in to Cloudera Manager Ranger Admin Web UI, as a Ranger administrator.
2. On Ranger Admin Web UI Service Manager Resource Knox, click cm\_knox.
3. In Knox Policies, open the CDP Proxy UI, API and Token policy.
4. In Knox Topology\*, add cdp-proxy-token.
5. Click Save.
6. Restart Ranger.

**CDPD-58704: hadoop roll key / key delete command shows operation failed error when one KMS host is down, even when operation succeeds**

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered kms instances. if one KMS instance is stopped (not deleted), client gets 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.**

```
hbase:001:0> revoke 'hrt_11'
ERROR: org.apache.hadoop.hbase.coprocessor.CoprocessorException: HTTP 400 Error: processSecureRevokeRequest processing failed
 at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocessor.preRevoke(RangerAuthorizationCoprocessor.java:1309)
 at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocessor.preRevoke(RangerAuthorizationCoprocessor.java:1128)
 at org.apache.hadoop.hbase.master.MasterCoprocessorHost$162.call(MasterCoprocessorHost.java:1857)
 at org.apache.hadoop.hbase.master.MasterCoprocessorHost$162.call(MasterCoprocessorHost.java:1854)
 at org.apache.hadoop.hbase.coprocessor.CoprocessorHost$0bserverOperationWithoutResult.callObserver(CoprocessorHost.java:558)
 at org.apache.hadoop.hbase.coprocessor.CoprocessorHost.executeOperation(CoprocessorHost.java:631)
 at org.apache.hadoop.hbase.master.MasterCoprocessorHost.preRevoke(MasterCoprocessorHost.java:1854)
 at org.apache.hadoop.hbase.master.MasterRpcServices.revoke(MasterRpcServices.java:2740)
 at org.apache.hadoop.hbase.shaded.protobuf.generated.MasterProtos$MasterService$2.callBlockingMethod(MasterProtos.java)
 at org.apache.hadoop.hbase.ipc.RpcServer.call(RpcServer.java:387)
 at org.apache.hadoop.hbase.ipc.CallRunner.run(CallRunner.java:139)
 at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:369)
 at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:349)
```

None

**CDPD-56741: Improvement in log message when jwtauth not used**

The following exception is printed at startup only and it is not cluttering logs:

```
2023-05-30 06:18:40,127 ERROR org.apache.ranger.rms.security.RMS
JwtAuthFilter:
 quasar-pibgzl-1.quasar-pibgzl.root.hwx.site-startStop-1]:
 Failed to initialize Ranger RMS JWT Auth Filter.
 java.lang.Exception: RangerJwtAuthHandler:
 Mandatory configs ('jwks.provider-url' & 'jwt.public-key')
 are missing, must provide atleast one.
 at org.apache.ranger.authz.handler.jwt.RangerJwtAuthHan
 dler.initialize(RangerJwtAuthHandler.java:84)
 ~[ranger-authn-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
 at org.apache.ranger.rms.security.RMSJwtAuthFilter.init
 ialize(RMSJwtAuthFilter.java:77)
 ~[ranger-rms-common-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0
 -186]
 at jdk.internal.reflect.NativeMethodAccessorImpl.invoke
 0(Native Method) ~[?:?]
 at jdk.internal.reflect.NativeMethodAccessorImpl.invoke(N
 ativeMethodAccessorImp
```

This exception is logged, since mandatory JWT auth filter configs are not provided (expected for some environments like y-cloud) Even though JWT auth filter is failed to initialised, it falls back to kerberos auth, hence no impact from auth perspective

NA

**CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup**

This is a warning log printed in catalina.out file when Ranger RMS server is initialised. Following exception is observed only in Oracle 19 setup. FileNotFoundException: /usr/share/java/oraclepki.jar

NA

**CDPD-48975: Ranger KMS KTS to KMS DB migration : keys with the same name but different case are not migrated**

KMS keys are not case sensitive

No workarounds. Such keys combination are very rare and migration doc was updated to check such keys before starting the migration.

**CDPD-42598: Kafka policy creation allowed with incorrect permissions.**

When creating a Kafka policy from the UI, the permissions "Idempotent write" and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource

None.

**CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload:**

```
{"resourceName": "classification", "userInput": "", "resources": {"classification": []}}
```

expectation is to return all the classifications . But the response has only "classification" Happens similarly for entity-label , entity-business-metadata.

None.

**CDPD-40734: User allowed to insert data into a hive table when there is a deny policy on a table column.**

A user is allowed to enter data into a table even if there is a deny policy present on one of the table columns.



Test scenario details:

```
Policy setup :-
 policy 1 :- all access policy for hrt_qa, hive and impala
 users
 resources - database - * , table - *, column - *
 users : hrt_qa, hive, impala
 access - all access allowed
 policy 2 :- policy on test_1.table_1 for hrt_5
 users : hrt_5
 resources : database - test_1, table - table_1, column - *
 access :- all access allowed
 policy 3 :- deny policy on test_1.table_1.c0 for hrt_5
 users : hrt_5
 resources : database - test_1, table - table_1, column -
c0
 access - all access denied
data setup :-
 database - test_1
 table - table_1(c0 int, c1 int)
```

The user is able to insert data into the table.

None.

### Known Issues in Spark

Learn about the known issues in Spark the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 9.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

The following known issues were identified in this release:

No new known issues in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9

#### CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the RuntimeWarning: Failed to add file message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

#### CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

**CDPD-67517: Spark3 tests are failing if /tmp is mounted as NOEXEC.**

Map the tmpdir to a writable path in spark3-conf/spark-defaults.conf using the following steps:

1. Go to the Cloudera Manager.
2. Go to the Spark service.
3. Click the Configuration tab.
4. Select Scope Gateway .
5. Select Category Advanced .
6. Locate the Spark Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-defaults.conf\_client\_config\_safety\_valve property.
7. Map the tmpdir to a writable path:

```
spark.driver.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
spark.executor.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
```

8. Enter a Reason for change, and then click Save Changes to commit the changes.
9. Deploy the client configuration.

**CDPD-23817: In the upgraded Cluster, the permission of /tmp/spark is restricted due to the HDP configuration hive.exec.scratchdir=/tmp/spark.**

If you are using the /tmp/spark directory in the CDP cluster, you must provide the required additional Policies/ACL permissions.

**CDPD-22670 and CDPD-23103: There are two configurations in Spark, "Atlas dependency" and "spark\_lineage\_enabled", which are conflicted. The issue is when Atlas dependency is turned off but spark\_lineage\_enabled is turned on.**

Run Spark application, Spark will log some error message and cannot continue. That can be restored by correcting the configurations and restarting Spark component with distributing client configurations.

**CDPD-23007: Mismatch in the Spark Default DB Location. In HDP 3.1.5, hive\_db entities have one attribute - 'location' which is configured to the '/managed' path. In fresh install of CDP 7.1.7, hive\_db entities now have 2 attributes 'location' configured to '/external' path and 'managedLocation' configured to '/managed' path. In the AM2CM migration (HDP 3.1.5 -> CDP 7.1.7), the 'location' attribute from hive\_db entities in HDP 3.1.5 comes unaltered to CDP 7.1.7 and hence maps to '/managed' path.**

This issue arises only if you are upgrading from HDP 3.1.5 to CDP 7.1.7. If you are performing a fresh install of CDP 7.1.7, you can ignore this issue.

None

**CDPD-217: The Apache Spark connector is not supported**

The old *Apache Spark - Apache HBase Connector* (shc) is not supported in CDP releases.

Use the new HBase-Spark connector shipped in CDP release.

**CDPD-3038: Launching pyspark displays several HiveConf warning messages**

When pyspark starts, several Hive configuration warning messages are displayed, similar to the following:

```
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of name hive.meta
store.runworker.in does not exist
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of
name hive.masking.algo does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of
name hive.metastore.runworker.in does not exist
```

```
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of na
me hive.masking.algo does not exist
```

These errors can be safely ignored.

### Known Issues in Schema Registry

Learn about the known issues in Schema Registry, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9

#### CDPD-40380: Authorization checking issue when Kerberos is disabled

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the `ranger.plugin.schema-registry.service.name` field.
3. Replace `GENERATED_RANGER_SERVICE_NAME` with the actual name of the service.
4. Restart the Schema Registry service.

#### CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

#### CDPD-59015: Schema Registry does not create new versions of schemas even if the schema is changed

Schema Registry uses a schema fingerprinting mechanism to differentiate between schemas. However, fingerprinting does not take into consideration the schema attributes of the field type. As a result, if you have two schemas where the only difference is that one has type attributes defined and the other does not, they will be considered identical by Schema Registry. For example, the following schemas are considered identical:

```
#Schema V1
{"type":"record","name":"schema_name","namespace":"ns","fields":
[{"name":"local_timestamp_micros_long","type":"long"}]}

#Schema V2
{"type":"record","name":"schema_name","namespace":"ns","fields":
[{"name":"local_timestamp_micros_long","type":{"type":"long","logicalType":"local-timestamp-micros"}}]}
```

Notice that the only difference is that in the second schema, the `local_timestamp_micros_long` field has a logical type specified. In cases like this, the new version of the schema is not created, the initial version is used. This is true even if the data that is being produced has a new schema version. The ID of the first schema version is used and is put in the serialized record. The new schema version is not created.

This issue is common when using change data capture (CDC) connectors like the Debezium connectors. This is because CDC connectors create schemas with the logical type decimal based on the column type in the database schema. For example:

```
{"type":"record","name":"schema_name","namespace":"ns","fields":
[{"name":"database_column","type":{"type":"bytes","logicalType":"decimal","precision":64,"scale":0}}]}
```

If the database schema changes (for example, the column type), it is possible that only scale changes, which is a schema attribute.

```
{ "type": "record", "name": "schema_name", "namespace": "ns", "fields":
 [{ "name": "database_column", "type": { "type": "bytes", "logicalType":
 "decimal", "precision": 64, "scale": 1 } }] }
```

In this case, even though scale changed to 1, the first version of the schema is used where scale is 0. As a result, the data is consumed with the wrong scale.

Avoid using logical types or other attributes. Alternatively, ensure that there are no changes in the logical types or other attributes between schema versions.

#### **OPSAPS-68708: Schema Registry might fail to start if a load balancer address is specified in Ranger**

Schema Registry does not start if the address specified in the Load Balancer Address Ranger property does not end with a trailing slash (/).

Set the value of the RANGER\_REST\_URL Schema Registry environment variable to an address that includes a trailing slash.

1. In Cloudera Manager, select the Schema Registry service.
2. Go to Configuration.
3. Find the Schema Registry Server Environment Advanced Configuration Snippet (Safety Valve) property and add the following:

```
Key: RANGER_REST_URL
Value: [***RANGER REST API URL***]
```

Replace [\*\*\*RANGER REST API URL\*\*\*] with an address that can be used by Schema Registry to access Ranger. Ensure that the address ends with a trailing slash. For example: <http://ranger-1.cloudera.com:6182/>

4. Restart the Schema Registry service.

#### **OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade**

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

1. Access the Ranger Console (Ranger Admin web UI).
2. Click the cm\_atlas resource-based service.
3. Add the schemaregistry user to the all - \* policies.
4. Click Manage Service Edit Service .
5. Add the schemaregistry user to the default.policy.users property.

#### **OPSAPS-69317: Kafka Connect Rolling Restart Check fails if SSL Client authentication is required**

The rolling restart action does not work in Kafka Connect when the ssl.client.auth option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set ssl.client.auth to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

#### **Known Issues in Sqoop**

Learn about the known issues in Sqoop, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 9.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9

#### CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version

If you have upgraded to CDP PvC Base 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to CDP PvC Base 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see [Troubleshooting Apache Sqoop issues](#).

#### CDPD-44431: Using direct mode causes problems

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the `--direct` option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the `sqoop.enable.deprecated.direct` property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through `-Dsqoop.enable.deprecated.direct=true`.

#### CDPD-3089: Avro, S3, and HCat do not work together properly

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

#### Parquet columns inadvertently renamed

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.

#### PARQUET-99: Importing Parquet files might cause out-of-memory (OOM) errors

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.

### 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 in Cloudera Runtime 7.1.9 SP1 CHF 8.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9

#### OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager

Cloudera Manager does not display a Log Files menu for SMM UI role (and SMM UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by SMM UI is not supported by Cloudera Manager.

View the SMM UI logs on the host.

#### OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL\_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in SMM safety valve in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Add the following value for bootstrap servers>Save Changes > Restart SMM.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

#### **CDPD-70313: KNOX does not send Authentication header on FIPS configuration**

KNOX does not send neither the Authentication header nor the hadoop.auth cookie. Because of this, the SMM UI returns an HTTP 401 response and sets the www-authenticate: Negotiate header. After this, KNOX still does not send the Authentication header. This results in the SMM UI being inaccessible through Knox.

You can access the SMM UI directly in Cloudera Manager at Clusters SMM Streams Messaging Manager Web UI and log in using the Kerberos username and password.

#### **CDPD-82560: Streams Messaging Manager known unprotected endpoint**

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**

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**

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.

#### **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 in Cloudera Runtime 7.1.9 SP1 CHF 8.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9**

There are no new known issues identified in this release.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9**

##### **CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic**

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

#### **CDPD-11079: Blacklisted topics appear in the list of replicated topics**

If a topic was originally replicated but was later excluded for replication, 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, this topic will be included in the response. Additionally, the excluded topic will also be visible in the SMM UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

#### **CDPD-30275: SRM may automatically re-create deleted topics on target clusters**

If auto.create.topics.enable is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with srm-control. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1]
```

2. Wait until SRM is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

#### **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 in Cloudera Runtime 7.1.9 SP1 CHF 9.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9**

##### **COMPX-22434: Queue Manager fails to recreate Dynamic Queue Schedules (DQS) rule after exiting the mixed mode**

QueueManager fails to recreate the DQS rule after exiting the Yarn mixed mode. This occurs when the cluster is upgraded from 7.1.9 GA+7.11.3 to 7.1.9 SP1 CHF9.

This error does not occur on a fresh install of 7.1.9 SP1 CHF 9, and also during 7.1.9 SP1 GA to 7.1.9 SP1 CHF9 upgrade.

None.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9**

##### **COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."**

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

None.

##### **COMPX-4644: Queue capacity rounding problem when configuration is initially set via YARN**

When setting the capacity scheduler configuration through the YARN/Cloudera Manager configuration, there may be capacity values that use multiple decimal places. This results in rounding/floating point precision discrepancies in the UI when trying to validate that all sibling capacities equal 100%. The UI looks like all the numbers add up to 100, but the validation still displays an error and does not allow to save the capacities. It is also observed that the capacity is being calculated as, for example, 99.9999999991 in the backend.

- Create queues within the UI, or
- Ensure that capacities configured through the Capacity Scheduler safety valve do not have more than one decimal place.

## 20202 Database migration after enabling opt-in migration

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to CDP 7.1.9, forcing migration from H2 to PostgreSQL database.
- Upgrade to CDP 7.1.9 CHF2, moving back to H2 database.
- Upgrade to CDP 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version CDP 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

## CDPD-56559: MapReduce jobs can intermittently fail during a rolling upgrade.

During a rolling upgrade between CDP versions 7.1.8 and 7.1.9, MapReduce jobs may fail with message, `RuntimeException: native snappy library not available`. Although the native Snappy compression library is not loaded, a checkmark displays to indicate that the Snappy compression library is loading for NodeManagers that are pending upgrades. This causes the MapReduce jobs that are associated with the NodeManagers to fail. After the upgrade, the jobs work as expected. This issue only impacts rolling upgrades from before CDP 7.1.9 to a higher version.

None.

## COMPX-12021 Queue Manager configurations on Scheduler Configuration page are not working

When setting the following properties on the YARN Queue Manager UI, the properties are set in the `capacity-scheduler.xml` file which does not have any effect on YARN. The properties need to be set in the `yarn-site.xml` file, which does not happen when you set them through YARN Queue Manager.

- "Maximum Application Priority" – `"yarn.cluster.max-application-priority"`
- "Enable Monitoring Policies" – `"yarn.resourcemanager.scheduler.monitor.enable"`
- "Monitoring Policies" – `"yarn.resourcemanager.scheduler.monitor.policies"`



**Note:** You can also set this property on the YARN configuration page in Cloudera Manager as "Capacity Scheduler Preemption".

- "Preemption: Observe Only" – `"yarn.resourcemanager.monitor.capacity.preemption.observe_only"`
- "Preemption: Monitoring Interval (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.monitoring_interval"`
- "Preemption: Maximum Wait Before Kill (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.max_wait_before_kill"`
- "Preemption: Total Resources Per Round" – `"yarn.resourcemanager.monitor.capacity.preemption.total_preemption_per_round"`
- "Preemption: Over Capacity Tolerance" – `"yarn.resourcemanager.monitor.capacity.preemption.max_ignored_over_capacity"`
- "Preemption: Maximum Termination Factor" – `"yarn.resourcemanager.monitor.capacity.preemption.natural_termination_factor"`
- "Enable Intra Queue Preemption" – `"yarn.resourcemanager.monitor.capacity.preemption.intra-queue-preemption.enabled"`

1. In Cloudera Manager, select the YARN service.
2. Click the Configuration tab.
3. Search for `yarn-site.xml`.
4. Under YARN Service Advanced Configuration Snippet (Safety Valve) for `yarn-site.xml`, add the corresponding parameter and value you need.
5. Click Save Changes.
6. Restart the YARN services.



### **COMPX-13177: QueueManager webapp requests fail with 'HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS\_EDH\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA'**

Products:

- Cloudera Manager for CDP Private
- Cloud Base Cloudera Manager for CDP Public Cloud

Context:

- Centos 7.8 and Redhat 7.8 operating systems, when FIPS support is enabled.

Problem:

- When attempting to display the Yarn Queue Manager interface, Cloudera Manager displays an error: "HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS\_EDH\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA".

1. Edit the file [/etc/default/cloudera-scm-server]
2. Around line 28, modify the line that starts with

```
#export CMF_OVERRIDE_TLS_CIPHERS=...
```

3. Remove the comment mark #.
4. Remove all ciphers with "3DES" in the name.
5. Save the file.
6. Restart the Cloudera Manager Server service.

### **Third-party applications do not launch if MapReduce framework path is not included in the client configuration**

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the `mapreduce.application.framework.path` property is set to the appropriate value, but third-party applications with their own configurations does not launch.

Set the `mapreduce.application.framework.path` property to the appropriate configuration for third-party applications.

### **JobHistory URL mismatch after server relocation**

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated the `mapred-site.xml` file that references the correct JobHistory Server.

### **YARN cannot start if Kerberos principal name is changed**

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN does not start. In such cases, the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the `znode` and restart the YARN service.
- Use the `reset ZK ACLs` command. This also sets the `znodes` below `/rmstore/ZKRMStateRoot` to `world:anyone:cdw` which is less secure.

### **Queue Manager does not open on using a custom user with a default Kerberos principal**

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager `process_username` property matches the YARN `process_username` property.

### Known Issues in Zeppelin

Learn about the known issues in Zeppelin the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 9.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9

#### 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.

#### BUG-125263: Zeppelin service move fails on clusters upgraded from HDP3.1.5

Resolve the circular symlink issue on the Zeppelin node by linking the conf directory to a new directory under /etc/zeppelin:

- # mkdir -p /etc/zeppelin/<version>/0
- # rm /usr/hdp/<version>/zeppelin/conf
- # ln -s /etc/zeppelin/<version>/0 /usr/hdp/<version>/zeppelin/conf

Where version is the HDP version. For example, 7.1.x-yyy. Restart the Zeppelin server in Ambari.

#### CDPD-3090: Due to a configuration typo, functionality involving notebook repositories does not work

Due to a missing closing brace, access to the notebook repositories API is blocked by default.

From the CDP Management Console, go to Cloudera Manager for the cluster running Zeppelin. On the Zeppelin configuration page (Zeppelin serviceConfiguration), enter shiro urls in the Search field, and then add the missing closing brace to the notebook-repositories URL, as follows:

```
/api/notebook-repositories/** = authc, roles[{{zeppelin_admin_group}}]
```

Click Save Changes, and restart the Zeppelin service.

#### CDPD-2406: Logout button does not work

Clicking the Logout button in the Zeppelin UI logs you out, but then immediately logs you back in using SSO.

Close the browser.

#### OPSAPS-59802: Zeppelin and Livy roles should be co-located on the same host.

When installing or upgrading to CDP Private Cloud Base, you must co-locate all Zeppelin and Livy roles on the same cluster host due to an issue with certificate generation.

### Known Issues in ZooKeeper

Learn about the known issues in ZooKeeper the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 9.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 9

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 9

There are no new known issues identified in the previous releases.

## Cumulative hotfix 8

Know more about the cumulative hotfix 8 for CDP 7.1.9 SP1. This cumulative hotfix was released on May 8, 2025.

**Cloudera Runtime 7.1.9.1042 (Cumulative Hotfix 8) download URL:****Parcel Repository Location**

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1042/parcels/
```

**What's new in Cloudera Runtime 7.1.9 SP1 CHF 8**

Understand the functionalities and improvements to features of components in Cloudera Runtime 7.1.9 SP1 CHF 8.

**New OS Support**

- RHEL 9.5
- RHEL 8.10 FIPS (extended RHEL 8.10 support for FIPS customers)

**What's new in Streams Replication Manager****Enhanced security with HTTP headers and HSTS enforcement**

Streams Replication Manager now includes various security headers in all API endpoint responses. These headers are added to improve overall security. Additionally, if SSL is enabled, HTTP Strict Transport Security (HSTS) is enforced as well to ensure that communication occurs over HTTPS.

The following headers are now included in API responses:

- Strict-Transport-Security: max-age=31536000; includeSubDomains - Enforces HTTPS communication for all requests, including subdomains, for a duration of one year. This header is only added when SSL is enabled.
- X-Frame-Options: SAMEORIGIN - Prevents clickjacking attacks by disallowing the page from being embedded in elements on other domains.
- X-XSS-Protection: 1; mode=block - Enables the XSS filter of the browser and blocks the page if an attack is detected.
- Content-Type: application/json - Ensures proper MIME type handling for responses.
- Cache-Control - Controls caching behavior.
  - Configured for API endpoints as no-store.
  - Configured for Swagger UI endpoints as max-age=7200.
- X-Content-Type-Options: nosniff - Prevents MIME type sniffing.
- Content-Security-Policy - Implements strict policies to control the sources of content that can be loaded.
  - Configured for API endpoints as default-src 'none';.
  - Configured for Swagger UI endpoints as default-src 'self' 'unsafe-inline'; font-src 'self' data:; img-src 'self' data:; style-src-attr 'self' 'unsafe-hashes'; style-src 'self' 'unsafe-inline'.

**Fixed issues in 7.1.9 SP1 CHF 8**

Know more about the cumulative hotfix 8 for 7.1.9 SP1.



**Note:** Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p1042.65851740.

**CDPD-78266: Ozone Manager displays NullPointerException (NPE) when overwriting empty file using multipart upload**

S3MultipartUploadCompleteRequestWithFSO displays NullPointerException when an empty file is being overwritten by non-zero file because omBucketInfo allows null value (null value is passed when non-update needed). This patch fix this by checking omBucketInfo before use.

**CDPD-82201: OMKeyAclRequestWithFSO is incorrectly setting full path as key name**

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.

**CDPD-64865: Intermittent timeout in TestBlockDeletion.testBlockDeletion ( BlockDeletionService stuck)**

If the event of SCM exiting safe mode was triggered multiple times, this might eventually cause the SCMBlockDeletingService to transition into a PAUSING state. This fix addresses the issue.

**CDPD-73278: Update OM, SCM, Datanode conf for RATIS-2135**

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-73736: DN Startup fails with "RuntimeException: Can't start the HDDS datanode plugin" error**

Remove the predefined hdds.ratis.raft.grpc.message.size. Its default value is determined by hdds.container.ratis.log.appender.queue.byte-limit + 1MB = 33MB.

Apache JIRA: [HDDS-11375](#)

**CDPD-78932: Container replication should be atomic**

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-80742: ConstraintViolationException was crashing the ContainerHealthTask in Ozone Recon**

ConstraintViolationException was crashing the ContainerHealthTask in Recon. After this fix, task will not crash and continue to identify the Unhealthy containers in SCM if any.

Apache JIRA: [HDDS-12585](#)

**CDPD-81939: Volume scanner should fail volume if rocksDB is inaccessible**

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-78384: Volume should not be marked as unhealthy when the disk is full**

When the disk is about to be full, ignore checking write check during volume scanner.

Apache JIRA: [HDDS-12239](#)

**CDPD-78960: Container import processing should respect reserved space at DN**

Reserve space during container import during replication similar to create container.

Apache JIRA: [HDDS-12235](#)

**CDPD-78506: Fix PATH environment variable creation for Shell action**

With this fix, in Oozie's shell actions, the PATH environment variable is evaluated based on the YARN NodeManager host's settings. So, the PATH environment variable applies to the Launcher AM container when executed, rather than being evaluated on the Oozie server's JVM (the legacy method). You can revert to the legacy behavior by setting oozie.action.shell.setup-path-in-oozie-server to true in oozie-site.xml.

Also, from now, Oozie also allows defining action-specific environment variables using oozie.launcher.<\*\*\*ACTION\_TYPE\*\*\*>.action.env.<\*\*\*VARIABLE\_NAME\*\*\*> allowing you to customize execution environments for each action type in the Launcher AM.

**CDPD-81435: Upgrade commons-vfs2 to 2.10 due to CVE-2025-27553**

The commons-vfs2 version is updated to 2.10 to fix vulnerability issues.

**CDPD-81401: Impala SQL queries that include the WITH clause should populate lineage in Atlas**

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-40346: The ddlQueries and ALTERNATIVE\_ADDCOLS lineage missing for Impala tables.**

The ALTERNATIVE\_ADDCOLS lineage no longer has issues when an Impala table is altered and the corresponding lineage is not created.

**CDPD-80921: Without a permission for one glossary, the /glossary call throws exception and it does not list the remaining glossaries**

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.

Apache JIRA: [ATLAS-4995](#)

**CDPD-78832: Livy:Bootstrap upgrade/replacement due to EOL and CVEs**

The bootstrap version is updated to the latest supported version.

**CDPD-77911: Missing Log4j Redactor dependency**

The class org.cloudera.log4j.redactor.RedactorAppender is made available in classpath after adding the org.cloudera.logredactor dependency to Atlas pom.xml. Providing the dependency ensures log redaction, avoiding potentially exposing credentials or PII in logs. This also prevents possible Log4j errors during startup.

**CDPD-61527: Accessing Impala service through Knox on FIPS clusters fails with SSL error code 5**

Resolved an issue that caused SSL error code 5 when accessing the Impala service through Knox on FIPS-enabled clusters.

**CDPD-81309: Ranger KMS with Oracle DB was not supported for Navigator Encrypt**

Navigator Encrypt deposit registration is failing with Ranger KMS DB with Oracle DB setup with the following error:

```
java.sql.SQLException: ORA-02289: sequence does not exist
Error Code: 2289
```

The issue is fixed now.

**CDPD-27801: Knox is missing HSTS header for HTTP 404 responses**

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. This feature introduces additional security layer for the HTTPS connections. For web applications using HSTS, browsers will be instructed to not perform attempts of unencrypted connections to particular domain even if they find URL-s for such.

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-81228: 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-77972: 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 by updating the `Heartbeater::threaf::last_tombstoned_report_time_` field. To re-enable the behavior, adjust the `--tserver_send_tombstoned_tablets_report_interval_secs` flag as needed.

**CDPD-80268: Hide logout button when configured with Knox**

In environments where Hue is configured with Knox, the Hue interface displayed 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-80574: Wrong results when CASE expressions has function calls referencing CHAR type expressions or columns**

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-78680: Selected long string values are getting truncated in the react-select control**

On the 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 the Ranger React UI.

**CDPD-81709: Update parquet-avro to 1.15.1 due to CVE-2025-30065**

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.

**CDPD-81755: Restrict trusted packages in the parquet-avro module**

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`.

**CDPD-81687: Performance improvement for DelegationTokenSecretManager**

This fix optimizes the object locking by KMS threads and improves the overall KMS throughput. The fix is a part of `hadoop-common`.

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

- [CVE-2024-1735](#) - Armeria

- [CVE-2025-27553](#) - Commons-vfs2
- [CVE-2025-30474](#) - Commons-vfs2
- [CVE-2025-27152](#) - Axios
- [CVE-2025-30065](#) - Apache Parquet

### Technical Service Bulletins

#### TSB-820: Potential Data Integrity Issues Found in Ozone

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-820: Potential Data Integrity Issues Found in Ozone](#)

### Known issues in 7.1.9 SP1 CHF 8

The following section lists the the known issues and limitations, the areas of impact, and workaround for each Runtime component.

#### Known Issues in Atlas

Learn about the known issues in Ozone the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8

The following known issues were identified in this release:

#### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

##### CDPD-77054: Exporting shell entities can cause the import to fail

For Apache Hive tables created with Apache Spark, there may be shell entities created in Apache Atlas. If these entities are in the export zip file before they are resolved, the import will fail.

Workaround: None

##### CDPD-70450: Impala SQL queries that include the “WITH” clause should populate lineage in Atlas

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-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

##### CDPD-6565: Issue with ddlQueries in Atlas for Hive/Impala tables created in default/non-default db using DAS/HUE respectively

ddlQueries are not created for Impala origin table. ddlQueries created for Impala/Hive CTAS tables have difference in names as compared to same operation for Hive table. The Impala table has the query in the ddlQueries entity name.

None

##### CDPD-55301: ddlQueries and ALTERNATE\_\* lineage missing for Spark tables created through spark3-shell

The ddlQueries and ALTERNATE\_\* lineage missing for Spark tables created using spark3-shell.

None

##### CDPD-56085: 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

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\*.

None

**CDPD-58554: Discard audits of specific classification, label, and business metadata**

Support to control audits for specific classification, label, and business metadata is not present in the Custom Audit Filters feature.

None

**CDPD-67112: Import transforms do not work as expected when replacing a string which already has ":"**

The character ":" is not supported in path replacements. The import succeeds but location remains unchanged. The character ":" must be avoided.

None

**CDPD-69150: Unable to add labels or user defined properties in Japanese**

Adding Japanese labels or user defined properties results in the error message: "Invalid label: ###, label should contain alphanumeric characters, \_ or -"

None

**CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties**

Entities, such as `hdfs_path`, are not returned by the search when their indexable properties are searched using partial search terms made of Japanese or Chinese characters. Only exact matches return results when searching the indexed properties made of Japanese or Chinese characters.

None

**CDPD-68191: Suggestions do not return the correct results when searching multiple Chinese characters**

Free text search does not return results when searching Chinese phrases made of multiple characters. Partial searches return the correct results.

None.

**CDPD-71219: Regression : Suggestions don't work for single character words on indexable attributes**

When searching for entities whose name (entity names are indexable) is a single character, search results are returned but suggestions are not. When searching for entities whose description (entity descriptions are not indexable) is a single character, both search results and suggestions are returned.

None.

**CDPD-67450: Table name renaming operation is not updating or creating iceberg\_table entity**

Renaming an Iceberg table does not update the corresponding Atlas entity.

None.

**CDPD-67089: Export/Import: When a table with Ozone path is exported as "connected", only the Ozone key is exported.**

When table with Ozone path is exported as "connected", only Ozone key is exported. Other Ozone entities, such as Ozone volume, Ozone bucket are not exported.

None.

**CDPD-43772: Performance issues with Atlas service**

If there are lot of update operations and the compression type of column families of `atlas_janus` table is SNAPPY, then the Kafka message processing might become slower.

- Consider setting compression type of column families of `atlas_janus` table as GZ.

**OPSAPS-67783: During rolling upgrade one among two Atlas server failed to start but Cloudera Manager considered as success**



Cloudera Manager marks the Execute command Start on service Atlas-1 as a success even when Atlas service had failed to start successfully. In such cases, Atlas logs give the exact reason for Atlas start-up failure.

None

**CDPD-19358: "IsIndexable" and "isOptional" value of a typedef's attribute is modified post migration.**

In HDP-265, falcon\_feed\_creation has "stored-in" attribute which has IsIndexable value set to True.

Post migration, "stored-in" is moved to relationshipAttributeDefs and has IsIndexable value set to False.

Similarly, in HDP-265 the table attribute of hive\_storage\_desc isOptional is set to true. Post migration, isOptional set to False.

None.

**CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command**

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**

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**

In converting content from Navigator to Atlas, the create time for Hive tables is not moved to Atlas.

None

**CDPD-10574: Suggestion order doesn't match search weights**

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**

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@58112
bc4 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) does not 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 the metadata collection from HMS. For example, if an ALTER TABLE operation in Spark that is changing a table name and is reporting 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 "@" does not fetch the correct results**

When searching Atlas qualifiedName values that include an "at" 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, and use 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 does 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 logout 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 CDP.

None.

**CDPD-1664: Guest users are redirected incorrectly**

Authenticated users logging in to Atlas are redirected to the CDP 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: The 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, to add metadata tags

to relationships that you wanted to make sure were unique in the system, you can design the relationship attribute with the property `IsUnique` equal `true`. However, in this release, the `IsUnique` attribute is not enforced.

None.

**DOCS-13759: Tag Propagation stops after a certain depth while the lineage is being extended**

When a tag is added to an entity at timestamp T1, the entities along the lineage to which the tag must be propagated is calculated at T1. Before tag propagation completes, if the lineage is extended, tag does not propagate to the entities in the extended lineage.

**CDPD-41142: When a Kafka console consumer group is run, more than one update audits are seen**

After running the console consumer with a consumer group, verify the consumer group entity created, along with the metrics and notifications for the consumer group and topic. The expected result can be: one `ENTITY_CREATE` audit and one `ENTITY_UPDATE` audit. But more than one `ENTITY_UPDATE` audits are seen.

**CDPD-40165: Two audits are created for SPARK CTAS table**

When following Spark queries are fired:

```
spark.sql("create table table1(id int)")
spark.sql("create table table2 as select * from table1")
```

HMS sends `"ENTITY_CREATE"` and `"ENTITY_FULL_UPDATE_V2"`.

The extra `ENTITY_FULL_UPDATE_V2` message received from HMS is sent as a part of `ALTER TABLE ADDCOLS` event from the HMS Hook side. This behaviour is observed only when the queries are run from Spark SQL and not when run the same queries are run from Beeline.

**CDPD-39197: Debug metrics returns empty data**

When debug metrics is enabled and some operations are performed, the response is empty

**CDPD-36495: Updating legacyAttribute from False to True resets the initially created relationshipAttributes values**

Creating types, entities, and to start, you must set the relationship with `is_legacy_attribute` value as **False**.

Later, update the value `relationshipDef is_legacy_attribute` to **True**.

For the entities that were created before updating the `is_legacy_attribute` to **True**, the `relationshipAttributes` value is reset.

**CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities**

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

**CDPD-22744: Bulk entity DELETE API does not delete authorised entities**

Bulk entity `DELETE` API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

**CDPD-29409: Hive import: Suggestion suggests entity which is deleted.**

Suggestions suggests tables of a database, which is a deleted entity.

**CDPD-25152: Tag propagation through deferred actions consumes additional time as compared to default flow**

The additional time might be due to the small overhead added to create / update task vertex and which is run in the background. This also depends on number of tasks queued to be executed.

**CDPD-42954: Zeppelin notebook fails after enabling Atlas-HDFS hook**

The Zeppelin notebooks are failing with errors after enabling Atlas-HDFS hook in the CDP cluster.

When the below properties are set for `atlas-client.properties` in Cloudera Manager:

- atlas.jaas.KafkaClient.option.keyTab
- atlas.jaas.KafkaClient.option.principal

Along with adding the properties in /etc/atlas/conf/atlas-application.properties, Cloudera Manager also adds these properties to atlas-application.properties for other services (like Spark).

Adding these properties interferes with the normal flow of the services (like Spark)

To enable HDFS lineage feature, instead of setting these properties through Cloudera Manager, users can manually add the properties directly in /etc/atlas/conf/atlas-application.properties

#### **CDPD-10576: Deleted Business Metadata attributes appear in Search Suggestions**

Atlas search suggestions continue to show Business Metadata attributes even if the attributes have been deleted.

None.

#### **CDPD-39427: [HDFS Lineage]: When the input is a directory in case of put/copyfromLocal/cp/mv commands, lineage is not created even though the script succeeds.**

When Source is a directory and target is a directory which is already present in Atlas, the command succeeds and inserts the data in the desired location, but lineage is not created.

#### **DOCS-13760: System Attributes search, \_\_classificationNames: Search with parent tag does not return entities associated to its children tags**

System attribute search with \_\_classificationNames = parent\_tag returns entities associated to parent\_tag only and not entities associated to its children tag.

Workaround: Instead of using system attribute, employ the basic search attribute "classification" which lists entities associated with inherited classifications.

#### **CDPD-35818: Basic search with tag filter provides approximateCount as -1 when there is no match and is 0 otherwise**

When the following search operations are performed:

- Faceted search with both tag filter and entity filter

The observed approximateCount is -1.

Here both entity filter and tag filter are present and when there is no match the response received is -1.

- Faceted search with only entity filter

Performing a basic search provides an approximate value of 0 when there is no match.

- Faceted search with only tag filter

Whenever there is tag filter in the query and there is no entity match, the approximateCount is -1 and if the tag filter is not available, the response approximateCount is 0

#### **CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities**

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

#### **CDPD-22744: Bulk entity DELETE API does not delete authorised entities**

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

#### **CDPD-76035: Resource lookup for Atlas service is failing**

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

#### **CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout**

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

**CDPD-79160: NPE while deleting BusinessMetadata**

If business metadata is created without adding any applicable types, a NullPointerException is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

**CDPD-82056: UI: when server response date fields as '0', UI shows as current time**

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](#)

**CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export**

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**

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](#)

**CDPD-66938: [Analyze] [Atlas] test\_time\_range tests fail**

When the Apache Atlas server is running on a node which has time zone other than UTC, there might be a time of day when the search results might differ if the relative CreateTime date range filters of TODAY, YESTERDAY, etc. are used. 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.

Use explicit date range filters instead of using relative date range filters, such as, TODAY, YESTERDAY.

**CDPD-84502: Advanced Search is not working properly in Atlas UI**

The Advance Search in Atlas does not return any results when selecting the Typedef option. The API is not triggered for each relation.

None

**CDPD-76789: Creating tag with name description throws java.lang.ClassCastException**

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails**

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

**CDPD-70321: Atlas Parallel import is failing with various errors**

During a parallel import-export activity with six iceberg table policies with exportOption as db1.\* for all six exports, all import fail after the exports.

**CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen**

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

**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**

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn**

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

**CDPD-79099: Unable to import hive tables using import script in Public Cloud**

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-88230: Bulk import operations currently do not construct lineage as expected**

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**

Bulk import operations do not create audits for the imported entities.

**Known Issues in Cruise Control**

Learn about the known issues in Cruise Control, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

**Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8**

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

#### Rebalancing with Cruise Control does not work due to the metric reporter failing to report the CPU usage metric

On the Kafka broker, the Cruise control metric reporter plugin may fail to report the CPU usage metric.

If the CPU usage metric is not reported, the numValidWindows in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.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 the [Moving a Host Between Clusters](#).



**Note:** Cluster nodes affected by this issue are not displayed as unhealthy.

### Known Issues in hive

Learn about the known issues in Hive, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

#### CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

**CDPD-60770: Beeline Authentication Issue with Special Characters in Passwords**

When LDAP is enabled, users cannot authenticate with Beeline if the password contains a special character. For example, the following string fails:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>;user=user@XXX;password='R3G#xpXyoy1M0Jb1'
```

Use the `-p` parameter to execute the Beeline command:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>; -n user@XXX -p 'R3G#xpXyoy1M0Jb1'
```

**CDPD-43769: Inconsistent behavior with CHAR comparisons and string literals**

In Hive, 'char' data to 'string' data comparison is considered as equal, but 'char' data to 'varchar' data comparison is considered as not equal if the data has trailing spaces. For example,

```
create table test(colchar char(10), colvarchar varchar(20), colstring string);
insert into test values ('a ', 'a ', 'a ');

select count(*) from test where colchar = colstring;
Output returns 1 because char and string data is considered as equal

select count(*) from test where colchar = colvarchar;
Output returns 0 because char and varchar data is considered as not equal
```

It is recommended that you use 'varchar or string' data type instead of 'char' data type for any character or string representation.

**CDPD-42726: Wrong results for agg queries from stats after running multi-insert query on managed table**

Multiple inserts at the same time into the same partition results in invalid stats. For example:

```
create table source(p int, key int, value string);
insert into source(p, key, value) values (101,42,'string42');

create table stats_part(key int,value string) partitioned by (p int);
from source
insert into stats_part select key, value, p
insert into stats_part select key, value, p;

select count(*) from stats_part;
-- In this scenario, StatsOptimizer helps serving this query because the result should be rowNum of the partition p=101.
-- The result is 1, however, the result should be 2.
```

Rewrite the multi-insert statement by merging the branches inserting into the same partition to one branch.

**CDPD-57574: Query execution fails due to NullPointerException in DagUtils.setupQuickStart**

Hive queries may fail at runtime with a `NullPointerException` while executing the Tez graph at `DagUtils.setupQuickStart` method. A part of the stack trace from the error is shown below:

```
ERROR : Failed to execute tez graph.
java.lang.NullPointerException: null at org.apache.hadoop.hive.ql.exec.tez.DagUtils.setupQuickStart(DagUtils.java:1724)
```



The error usually occurs from an invalid plan that has cycles created by semi-join or map-join edges.

Disable some of the optimizations in order to remove the cycles from the plan. It is sufficient to set one of the following properties:

- set hive.tez.dynamic.partition.pruning=false;
- set hive.tez.dynamic.semijoin.reduction=false;
- set hive.auto.convert.join=false;

To avoid any undesired performance regressions, disable the optimization(s) only for specific queries and not globally (through Cloudera Manager) for the whole workload or cluster.

#### **CDPD-45370: Cartesian join fails with an empty table**

When the configuration hive.tez.cartesian-product.enabled is set to true, query fails if one of the tables involved in a Cartesian join has zero rows. This results in an IllegalArgumentException with the message Number of items is 0. Should be positive.

To address this issue, set hive.tez.cartesian-product.enabled to false before running the query.

Apache Jira: [\[HIVE-26582\]](#)

#### **CDPD-43107: SemanticException for INSERT INTO statement when Hive Cost-based Optimizer (CBO) is disabled**

If you are running the INSERT INTO query and hive.cbo.enable is set to "false", the query fails with a SemanticException. For example,

```
set hive.cbo.enable=false;

CREATE TABLE mytable (
 id INT,
 str STRING
);

INSERT INTO mytable (id, str) VALUES (1, 'a');
```

```
Output:
org.apache.hadoop.hive.ql.parse.SemanticException: 0:0 Expected
2 columns for insclause-0/default@mytable; select produces 1 col
umns. Error encountered near token ''a''
```

This issue occurs only when columns are specified for the target table in the query. INSERT INTO mytable values (1, 'a'); does not result in an exception.

Enable CBO and run the query.

#### **CDPD-43957: HiveServer shuts down during replication due to high resource usage**

During Hive replication (both bootstrap and incremental), you may notice that the HiveServer (HS2) shuts down periodically with the following error:

```
java.sql.SQLException: org.apache.hive.jdbc.ZooKeeperHiveClientE
xception: Unable to read HiveServer2 configs from ZooKeeper
 at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.ja
va:265)
 at org.apache.hive.jdbc.HiveDriver.connect(HiveDriver.java:107)
 at java.sql.DriverManager.getConnection(DriverManager.java:664)
 at java.sql.DriverManager.getConnection(DriverManager.java:247)
 at com.cloudera.enterprise.hive3qt.Hive3QueryTool$HiveOperation.
execute(Hive3QueryTool.java:682)
 at com.cloudera.enterprise.hive3qt.Hive3QueryTool.main(Hive3Qu
eryTool.java:935)
```

```

Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException:
Unable to read HiveServer2 configs from ZooKeeper
 at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:177)
 at org.apache.hive.jdbc.Utills.configureConnParamsFromZooKeeper(Utills.java:580)
 at org.apache.hive.jdbc.Utills.parseURL(Utills.java:391)
 at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.java:263)
 ... 5 more
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Tried all existing HiveServer2 uris from ZooKeeper.
 at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.getServerHosts(ZooKeeperHiveClientHelper.java:132)
 at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:172)
 ... 8 more

```

This issue occurs when you replicate at scale with an unbalanced cluster setup that has all the roles running on the Cloudera Manager host. As a result, Cloudera Manager ends up in a bottleneck situation because HS2 crashes, preventing further replication. The logs (/var/log/messages) indicate that there were a large number of Java processes running on the host, which exhausted the host's memory and triggered an Out Of Memory Killer process to stop HS2.

You can restart HS2 and reinitiate replication. However, the replication process may take a longer time to complete.

- Follow these recommendations when setting up your source and target clusters:
  - Ensure that no HS2, Hive metastore (HMS), DataNode, or NameNode roles are running on the Cloudera Manager host.
  - Ensure that you have multiple instances of HS2 and HMS roles on different nodes.
  - It is recommended that you have the HS2 and HMS roles on different nodes than the DataNodes or NameNodes.

These practices increase the possibility of Hive replication at scale completing successfully.

- Limit the scale of Hive replication by temporarily disabling replication policies.

#### CDPD-40730: Parquet change can cause incompatibility

Parquet files written by the parquet-mr library in CDP 7.1.9, where the schema contains a timestamp with no UTC conversion will not be compatible with older versions of Parquet readers. The effect is that the older versions will still consider these timestamps as they would require UTC conversions and will thus end up with a wrong result. You can encounter this problem only when you write Parquet-based tables using Hive, and tables have the non-default configuration `hive.parquet.write.int64.timestamp=true`.

None.

#### CDPD-41274: HWC + Oozie issue: Could not open client transport with JDBC Uri

Currently only Spark cluster mode is supported in the Oozie Spark Action with Hive Warehouse Connector (HWC).

Use Spark action in cluster mode.

```

<spark xmlns="uri:oozie:spark-action:1.0">
 ...
 <mode>cluster</mode>
 ...
</spark>

```

#### CDPD-26556 After an upgrade, querying a CTAS table under certain conditions might throw an exception

If you upgrade your Hive cluster from CDH 6 to CDP 7, create a CTAS table in the CDP cluster from a table you upgraded from CDH, you might see the following exception when you query the new table:

```
class org.apache.hadoop.io.IntWritable cannot be cast to class org.apache.hadoop.hive.serde2.objectinspector.StandardUnionObjectInspector$StandardUnion
```

This issue involves CDH-based tables having columns of complex types ARRAY, MAP, and STRUCT.

#### **CDPD-23506: OutOfMemoryError in LLAP**

Long running spark-shell applications can leave sessions in interactive Hiveserver2 until the Spark application finishes (user exists from spark-shell), causing memory pressure in case of a high number of queries in the same shell (1000+).

You must close spark-shell so that sessions are closed. Add the owner of the database or the tables as a user with read or read/write access to the tables directly.

#### **CDPD-23041: DROP TABLE on a table having an index does not work**

If you migrate a Hive table to CDP having an index, DROP TABLE does not drop the table. Hive no longer supports indexes ([HIVE-18448](#)). A foreign key constraint on the indexed table prevents dropping the table. Attempting to drop such a table results in the following error:

```
java.sql.BatchUpdateException: Cannot delete or update a parent row: a foreign key constraint fails ("hive"."IDXS", CONSTRAINT "IDXS_FK1" FOREIGN KEY ("ORIG_TBL_ID") REFERENCES "TBLS" ("TBL_ID"))
```

There are two workarounds:

- Drop the foreign key "IDXS\_FK1" on the "IDXS" table within the metastore. You can also manually drop indexes, but do not cascade any drops because the IDXS table includes references to "TBLS".
- Launch an older version of Hive, such as Hive 2.3 that includes IDXS in the DDL, and then drop the indexes as described in [Language Manual Indexing](#).

Apache Issue: [Hive-24815](#)

#### **CDPD-17766: Queries fail when using spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive.**

Do not specify spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive remotely.

Workaround: specify principal=hive.server2.authentication.kerberos.principal as shown in the following syntax:

```
jdbc:hive://<host>:<port>/<dbName>;principal=hive.server2.authentication.kerberos.principal;<otherSessionConfs>?<hiveConfs>#<hiveVars>
```

#### **CDPD-13636: Hive job fails with OutOfMemory exception in the Azure DE cluster**

Set the parameter hive.optimize.sort.dynamic.partition.threshold=0. Add this parameter in Cloudera Manager (Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml)

#### **CDPD-10848: HiveServer Web UI displays incorrect data**

If you enabled auto-TLS for TLS encryption, the HiveServer2 Web UI does not display the correct data in the following tables: Active Sessions, Open Queries, Last Max n Closed Queries

### Known Issues in Hue

Learn about the known issues in Hue, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

#### **CDPD-58978: Batch query execution using Hue fails with Kerberos error**

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark queries in a batch mode launches a shell script Oozie job from Hue and this is not supported on a secure cluster.

There is no workaround. You can submit the queries individually.

#### **OPSAPS-69659: Hue service fails on restart with "Unable to find pycogp2 2.5.4" error**

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.

1. Log in to Cloudera Manager as an Administrator.
2. Go to Clusters Hue Configurations and add the following key and value in the Hue Service Environment Advanced Configuration Snippet (Safety Valve) field:

Key: PYTHONPATH

Value: `[***PYTHON-PATH***]`

Replace `[***PYTHON-PATH***]` with the actual location where you have installed Python. For example, `/opt/cloudera/parcels/CDH/lib/hue/build/env/lib/python3.8/site-packages`

3. Click Save Changes.
4. Restart the Hue service.

#### **CDPD-59677: Unable to view Phoenix tables on the left assist in Hue**

On clusters secured with Knox, you may not be able to see Phoenix tables on the left assist that are present under the default database (that is, an empty("") database).

None.

#### **CDPD-58142: A query is not pre-populated in the Hue editor after clicking on the "Re Execute" button**

When you click Re Execute to rerun a query from the Job Browser Queries Query Details page, the query does not get populated on the Hue editor, as expected.

None.

#### **CDPD-54376: Clicking the home button on the File Browser page redirects to HDFS user directory**

When you are previewing a file on any supported filesystem, such as S3 or ABFS, and you click on the Home button, you are redirected to the HDFS user home directory instead of the user home directory on the said filesystem.

None.

#### **CDPD-41306: pip3.8 freeze command does not work and results into an error**

You see the `/usr/lib/hue/build/env/bin/python: No such file or directory` error when you run the following command:

```
build/env/bin/pip3.8 freeze
```

Run the freeze command as follows by specifying the paths of python3.8 and pip3.8:

```
/opt/cloudera/parcels/CDH/lib/hue/build/env/bin/python3.8 /opt/cloudera/parcels/CDH/lib/hue/build/env/bin/pip3.8 freeze
```

#### **CDPD-43293: Unable to import Impala table using Importer**

Creating Impala tables using the Hue Importer may fail.

If you have both Hive and Impala services installed on your cluster, then you can import the table using by selecting the Hive dialect from **Tables Sources**.

If only Impala service is installed on your cluster, then go to **Cloudera Manager Clusters Hue Configurations** and add the following line in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field:

```
[beeswax]
max_number_of_sessions=1
```

#### **DWX-8602: Unable to import a large CSV file from the local workstation**

You may see an error message while importing a CSV file into Hue from your workstation, stating that you cannot import files of size more than 200 KB.

Upload the file to S3 or ABFS and then import them into Hue using the Importer.

#### **CDPD-39330: Unable to use the pip command in CDP**

You may not be able to use the pip command in CDP 7.1.7 or higher and may see the following error when using pip in a command: “ImportError: cannot import name chardet”.

Follow the steps listed on [Unable to use pip command in CDP](#).

#### **CDPD-24294: Hue uses the unsafe-inline directive in its Content Security Policy (CSP) header**

Hue 4 web interface uses the unsafe-inline directive in its CSP header. As a result, the application server does not set the CSP header in its HTTP responses, and therefore does not benefit from the additional protection against potential cross-site scripting issues and other modern application vulnerabilities which a properly configured CSP may provide. This could lead to application vulnerability.

Cloudera recommends deploying additional security measures such as a firewall within the Hue server to control allowed connections, and SSO-based authentications mechanisms such as LDAP or SAML.

#### **OPSAPS-58927: Connection failed error when accessing the Search app (Solr) from Hue**

If you are using Solr with Hue to generate interactive dashboards and for indexing data, and if you have deployed two Solr services on your cluster and selected the second one as a dependency for Hue, then Cloudera Manager assigns the hostname of the first Solr service and the port number of the second Solr service generating an incorrect Solr URL in the search section of the `hue.ini` file. As a result, you may see a “Connection failed” error when you try to access the Search app from the Hue web UI.

1. Log into Cloudera Manager as an Administrator.
2. Go to **Clusters Hue service Configuration** and add the following lines in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field:

```
[search]
URL of the Solr Server
```

```
solr_url=http://[***HOSTNAME***]:[***PORT***]/solr/
```

For example:

```
solr_url=http://solr2:4567/solr/
```

3. Click Save Changes.
4. Restart the Hue service.

### CDPD-41136: Importing files from the local workstation is disabled by default

Cloudera has disabled the functionality to import files from your local workstation into Hue because it may cause errors. You may not see the Local File option in the Type drop-down menu on the **Importer** page by default.

You can enable the functionality to import files from your local workstation by specifying the following parameter in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue\_safety\_valve.ini field in Hue configurations in Cloudera Manager:

```
[indexer]
enable_direct_upload=true
```

### Invalid S3 URI error while accessing S3 bucket

The Hue Load Balancer merges the double slashes (//) in the S3 URI into a single slash (/) so that the URI prefix "/filebrowser/view=S3A://" is changed to "/filebrowser/view=S3A:". This results in an error when you try to access the S3 buckets from the Hue File Browser through the port 8889.

The Hue web UI displays the following error: “Unknown error occurred”.

The Hue server logs record the “ValueError: Invalid S3 URI: S3A” error.

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.

You should be able to load the S3 browser from both 8888 and 8889 ports.

Alternatively, you can use the Hue server port 8888 instead of the load balancer port 8889 to resolve this issue.

### 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.

- Restart the Hue Load Balancer.

#### **CDPD-17956, CDPD-14167: Setting idle session timeout for Hue does not work when the cluster is secured using Knox SSO**

If Hue is configured with `desktop.auth.backend.KnoxSpnegoDjangoBackend` as the Authentication Backend, then the automatic idle session logout that is set by configuring the `idle_session_timeout` property does not take effect. You may also see 404 error while accessing Hue from the Knox UI when the `idle_session_timeout` property is not set to -1.

None

#### **Query history displays "Result Expired" message**

You see the "Result Expired" message under the Query History column on the **Queries** tab for queries which were run back to back. This is a known behaviour.

None.

#### **CDPD-43956: Manually replace UUID when importing Oozie workflows containing sub-workflows**

When importing Oozie workflows that contain sub-workflows, you must replace all the UUID entries with unique new entries. If you change the UUID of a sub-workflow, you must update that reference in the parent workflow to avoid circular dependencies.

#### **Hue limitation after upgrading from CDH to CDP Private Cloud Base**

The `hive.server2.parallel.ops.in.session` configuration property changes from TRUE to FALSE after upgrading from CDH to CDP Private Cloud Base. Current versions of Hue are compatible with this property change; however, if you still would like to use an earlier version of Hue that was not compatible with this property being FALSE and shared a single JDBC connection to issue queries concurrently, the connection will no longer work after upgrading.

#### **Hue UI is blank upon login after upgrading to CDP 7.1.7 from CDH 6**

If your cluster was secured using Knox, and if you have upgraded from CDH 6 to CDP 7.1.7, then you may see a blank Hue screen. This could happen because the `knox_proxyhosts` parameter is newly introduced in CDP, and it is possible that this parameter is not configured in Cloudera Manager under Hue configuration.

Specify the host on which you have installed Knox in the Hue Knox Proxy Hosts configuration as follows:

- Log in to Cloudera Manager as an Administrator.
- Obtain the host name of the Knox Gateway by going to **Clusters Knox service Instances**.
- Go to **Clusters Hue service Configuration** and search for the Knox Proxy Hosts field.



**Note:** Cloudera Manager displays the following warning if the Knox Proxy Hosts field is empty when Knox Gateway is enabled on the CDP cluster: The parameter `knox_proxyhosts` cannot be empty. This can happen if there are no Knox Gateways. Please set the `knox_proxyhosts` to the list of hosts that have Knox Gateways.

- Specify the Knox Gateway hostname in the Knox Proxy Hosts field.
- Click **Save Changes** and restart the Hue service.

### **Unsupported features**

#### **CDPD-59595: Spark SQL does not work with all Livy servers that are configured for High Availability**

SparkSQL support in Hue with Livy servers in HA mode is not supported. Hue does not automatically connect to one of the Livy servers. You must specify the Livy server in the Hue Advanced Configuration Snippet as follows:

```
[desktop]
[spark]
```

```
livy_server_url=http(s)://[***LIVY-FOR-SPARK3-SERVER-HOST***]:
[***LIVY-FOR-SPARK3-SERVER-PORT***]
```

Moreover, you may see the following error in Hue when you submit a SparkSQL query: Expecting value: line 2 column 1 (char 1). This happens when the Livy server does not respond to the request from Hue.

Specify all different Livy servers in the `livy_server_url` property one at a time and use the one which does not cause the issue.

#### **CDPD-18491: PySpark and SparkSQL are not supported with Livy in Hue**

Hue does not support configuring and using PySpark and SparkSQL with Livy in CDP Private Cloud Base.

#### **Importing and exporting Oozie workflows across clusters and between different CDH versions is not supported**

You can export Oozie workflows, schedules, and bundles from Hue and import them only within the same cluster if the cluster is unchanged. You can migrate bundle and coordinator jobs with their workflows only if their arguments have not changed between the old and the new cluster. For example, hostnames, NameNode, Resource Manager names, YARN queue names, and all the other parameters defined in the `workflow.xml` and `job.properties` files.

Using the import-export feature to migrate data between clusters is not recommended. To migrate data between different versions of CDH, for example, from CDH 5 to CDP 7, you must take the dump of the Hue database on the old cluster, restore it on the new cluster, and set up the database in the new environment. Also, the authentication method on the old and the new cluster should be the same because the Oozie workflows are tied to a user ID, and the exact user ID needs to be present in the new environment so that when a user logs into Hue, they can access their respective workflows.



**Note:** Migrating Oozie workflows from HDP clusters is not supported.

#### **Known Issues in Kafka**

Learn about the known issues in Kafka, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8**

There are no new known issues identified in this release.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8**

##### **CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress**

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.



None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

#### **OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners**

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL\_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker) in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Save Changes > Restart SMM.

#### **The offsets.topic.replication.factor property must be less than or equal to the number of live brokers**

The offsets.topic.replication.factor broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a GROUP\_COORDINATOR\_NOT\_AVAILABLE error until the cluster size meets this replication factor requirement.

None

#### **Requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true**

The first few produce requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true.

Increase the number of retries in the producer configuration setting retries.

#### **Performance degradation when SSL Is enabled**

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with ssl.secure.random.implementation = SHA1PRNG. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

#### **OPSAPS-43236: Kafka garbage collection logs are written to the process directory**

By default Kafka garbage collection logs are written to the agent process directory. Changing the default path for these log files is currently unsupported.

None

#### **RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure**

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transactional method because we are in an error state
 at org.apache.kafka.clients.producer.internals.TransactionManager.maybeFailWithError(TransactionManager.java:1125)
 at org.apache.kafka.clients.producer.internals.TransactionManager.maybeAddPartition(TransactionManager.java:442)
 at org.apache.kafka.clients.producer.KafkaProducer.doSend(KafkaProducer.java:1000)
 at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:914)
 at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:800)
 .
 .
```

```
Caused by: org.apache.kafka.common.errors.Cluster
AuthorizationException: Cluster authorization failed.
```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting `enable.idempotence` to `false`.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

#### **CDPD-49304: AvroConverter does not support composite default values**

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

#### **DBZ-4990: The Debezium Db2 Source connector does not support schema evolution**

The Debezium Db2 Source connector does not support the evolution (updates) of schemas. In addition, schema change events are not emitted to the schema change topic if there is a change in the schema of a table that is in capture mode. For more information, see [DBZ-4990](#).

None.

#### **CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression**

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy compression. The HDFS Stateless Sink connector is only affected if the `Compression Codec` or `Compression Codec` for Parquet properties are set to `SNAPPY`.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: 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/na
tivelibs
```

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**

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 Knox**

Learn about the known issues in Knox, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8**

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

There are no known issues identified in earlier releases.

#### Known Issues in Kudu

Learn about the known issues in Kudu, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

**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.

Workaround: Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

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.

Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

### Unsupported features

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 in Cloudera Runtime 7.1.9 SP1 CHF 8.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8

The following known issues were identified in this release:

No new known issues in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

No known issues identified in Livy.

#### Known Issues in Iceberg

Learn about the known issues in Iceberg, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

There are no known issues identified in earlier releases.

### Known Issues in Impala

Learn about the known issues in Impala, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8

The following known issues were identified in this release:

#### CDPD-81980: Impala query may fail due to backpressure

Impala queries could fail with a dropped due to backpressure error, even when service queue and memory usage appear normal.

Increase the value of the `--mem_limit` startup option.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

#### CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

#### CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the `--tls_ciphersuites` startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.

Ensure that the list in the `--tls_ciphersuites` startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

**Apache Jira:** [IMPALA-14625](#)

#### CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

#### CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work

If you try to run `impala-shell` on RHEL 9 by setting the default python executable available in `PATH` to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, `IMPALA_PYTHON_EXECUTABLE=python3`.

#### CDPD-42958: After upgrading the CDH 7.1.9 from CDH 6.x, under certain conditions you cannot insert data into a table

Under the following conditions, after upgrading from CDH 6.x to CDH 7.1.9 you cannot insert data into a table from Impala:

- On CDH 6.x, you created a database with Impala in a user specified HDFS location.
- Using Hive, you then created a table in the database.

Under these conditions, the database and table are stored in the user-specified HDFS directory. After upgrading, the HDFS directory of the table is read-only for Impala. Consequently, from Impala you cannot insert new data into the table because Impala does not have write permission on the HDFS directory.

Workaround: To resolve this issue, use either one of the following workarounds:

- Using the Ranger Web UI, in the policy repository `cm_hdfs`, grant the user 'impala' write permission on the directory where the table resides.
- Enter the following command to grant write permission to user 'impala' on the HDFS directory where the table resides.

```
hdfs dfs -setfacl -m default:user:impala:rwX <HDFS directory>
```

### CDPD-80019: Impala fails to start due to large catalog-update topic

Impala coordinators fail to start if the catalog-update statestore topic exceeds 2GB, hitting the Thrift 0.16.0 message size limit (`thrift_rpc_max_message_size`). This prevents the coordinator from processing the update.

- Enable local catalog mode: `--use_local_catalog=true`
- In Cloudera Manager, go to Impala Configuration search for "Catalog Server Command Line Argument Advanced Configuration Snippet (Safety Valve)" and click 
- Add property name and value:

```
Name : --catalog_topic_mode
Value: minimal
```

- Save changes.

Apache Jira: [IMPALA-12350](#)

### 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/<data
base>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

**CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.**

You must use a single Knox node to access Impala UI.

**Impala api calls via 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-21828: Multiple permission assignment through grant is not working**

None

**Problem configuring masking on tables using Ranger**

The following Knowledge Base article describes the behavior when we configure masking on tables using Ranger. This configuration works for Hive, but breaks queries in some scenarios for Impala.

For a workaround, see the following Knowledge Base article: [ERROR: "AnalysisException: No matching function with signature: mask\(FLOAT\)" when Impala jobs fail with the following error with signature: mask\(FLOAT\)](#)

**IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS**

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm\_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXXXX"

**IMPALA-532: Impala should tolerate bad locale settings**

If the LC\_\* environment variables specify an unsupported locale, Impala does not start.

Add LC\_ALL="C" to the environment settings for both the Impala daemon and the Statestore daemon.

**Avro Scanner fails to parse some schemas**

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

**IMPALA-691: Process mem limit does not account for the JVM's memory usage**

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the impalad daemon.

To monitor overall memory usage, use the top command, or add the memory figures in the Impala web UI /memz tab to JVM memory usage shown on the /metrics tab.

**IMPALA-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon**

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Remove trailing semicolon from the Avro schema.

**IMPALA-1652: Incorrect results with basic predicate on CHAR typed column**

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the `RPAD()` function to blank-pad literals compared with `CHAR` columns to the expected length.

### IMPALA-1821: Casting scenarios with invalid/inconsistent results

Using a CAST() function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

### IMPALA-2005: A failed CTAS does not drop the table if the insert fails

If a `CREATE TABLE AS SELECT` operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

### Drop the new table manually after a failed CREATE TABLE AS SELECT

### IMPALA-2422: % escaping does not work correctly when occurs at the end in a LIKE clause

If the final character in the RHS argument of a LIKE operator is an escaped \% character, it does not match a % final character of the LHS argument.

None

## IMPALA-2603: Crash: impala::Coordinator::ValidateCollectionSlots

A query could encounter a serious error if includes multiple nested levels of INNER JOIN clauses involving subqueries.

None

### IMPALA-3094: Incorrect result due to constant evaluation in query with outer join

An OUTER JOIN query could omit some expected result rows due to a constant such as FALSE in another join clause. For example:

```

explain SELECT 1 FROM alltypestest a1
 INNER JOIN alltypesagg a2 ON a1.smallint_col = a2.year AND fals
e
 RIGHT JOIN alltypes a3 ON a1.year = a1.bigint_col;
+---+
| Explain String |
+---+
| Estimated Per-Host Requirements: Memory=1.00KB VCores=1 |
| 00:EMPTYSET |
+---+

```

### IMPALA-3509: Breakpad minidumps can be very large when the thread count is high

The size of the breakpad minidump files grows linearly with the number of threads. By default, each thread adds 8 KB to the minidump size. Minidump files could consume significant disk space when the daemons have a high number of threads.

Add `-minidump_size_limit_hint_kb=size` to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then



2 KB per thread after that.) The minidump file can still grow larger than the "hinted" size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

#### **IMPALA-4978: Impala requires FQDN from hostname command on Kerberized clusters**

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the `gethostname()` system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a Kerberized cluster.

Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname` only returns the short name, pass the command-line flag `##hostname=FULLY_QUALIFIED_DOMAIN_NAME` in the startup options of all Impala-related daemons.

#### **IMPALA-6671: Metadata operations block read-only operations on unrelated tables**

Metadata operations that change the state of a table, like `COMPUTE STATS` or `ALTER RECOVER PARTITIONS`, may delay metadata propagation of unrelated unloaded tables triggered by statements like `DESCRIBE` or `SELECT` queries.

None

#### **IMPALA-7072: Impala does not support Heimdal Kerberos**

None

#### **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.

#### **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.

#### **Known issues in Navigator Encrypt**

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround.

#### **Failed to start `navencrypt-mount.service`**

After installing NavEncrypt on SLES, if the command `"systemctl status navencrypt-mount"` fails with the error: "Failed to start `navencrypt-mount.service`: Unit `navencrypt-mount.service` failed to load: No such file or directory", the `systemd` files need to be reloaded.

Run the command:

```
(sudo) systemctl daemon-reload
```

#### **Crash in Navigator Encrypt kernel module if `--profile-file=` option is used with ACLs**

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 Ozone

Learn about the known issues in Ozone, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

#### CDPD-84518: Expired secret key can abort leader Ozone Manager startup

Ozone Manager fails to startup due to NPE exception because of expired SCM managed secret key.

Apache Jira: [HDDS-13234](#)

#### CDPD-81569: Recon is not reflecting correct metrics for Volumes, Buckets, Keys, and Pending Key Deletions

Recon is displaying data for container and space utilisation only. The metrics for Volumes, Buckets, Keys, and Pending Key Deletions are not being updated and displayed correctly.

Apache Jira: [HDDS-12607](#)

#### 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: 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.

#### The Prometheus binary is not available in CDP Private Cloud Base 7.1.9 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.

#### After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and if Ozone is in the Non-HA environment, an exception message is observed during the finalization of the Ozone upgrade.

During the finalization of the upgrade, a `ClassNotFoundException` message for `org.cloudera.log4j.redactor.RedactorAppender` class may be displayed. The error message can be ignored because the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

**CDPD-68951: In CDP Private Cloud Base 7.1.9 CHF2 version and lower, the command `ozone sh key list <bucket_path>` displays the `isFile` flag in a key's metadata as `false` even when the key is a file. This issue is rectified in CDP Private Cloud Base 7.1.9 CHF3. However, the pre-existing (pre-upgrade) key's metadata cannot be changed.**

None

**CDPD-67771: When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.**

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(OzoneManager
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.lambda
$1(OzoneManagerStateMachine.java:363) at
java.base/java.util.concurrent.CompletableFuture
$AsyncSupply.run(CompletableFuture.java:1700) at java.base/
java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1128)
at java.base/java.util.concurrent.ThreadPoolExecutor
$Worker.run(ThreadPoolExecutor.java:628) at java.base/
java.lang.Thread.run(Thread.java:834)
```

This occurs because S3A uses multipart upload to commit job results in a batch. The staging committer's replace mode deletes the target directory before completing Memory Protection Unit (MPU). The problem is that File System Optimization (FSO) does not create intermediate directories during MPU, it does only for regular file/dir/key requests.

Use `fs.s3a.committer.name=magic` for affected versions.

**CDPD-64398: In SCM with non-HA configuration, Secret key manager does not initialise. Hence, the startup of OM and DataNode fails because it cannot get a secret key. This key is used when security is enabled and used in Block token and container token verification while communicating between the Ozone client, OM, and Datanode.**

This issue does not occur in SCM with HA configuration.

You must force exit from the safe mode for SCM. This triggers the initialization of the Secret key manager. Below are the options:

- exit safe mode manually whenever SCM is started
- disable safe mode by setting `hdds.scm.safemode.enabled=false` in safety valve for SCM configuration

Impact of disabling safe mode and the purpose of safe mode:

- Write should not fail once the cluster is out of safe mode
- Read of existing data should not fail after the cluster is out of safe mode
- Unnecessary re-replication should be avoided during cluster restart

So disabling safe mode does not have any major impact on Ozone.

**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.

**CDPD-52412: keyManagerImpl#listStatus exceeds the maximum RPC length.**

To reduce the payload, you must increase the part size.

**OPSAPS-68159: If you did not deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Error when distributing parcel to host, Ignoring non-compliant parcel manifest error is displayed after Cloudera Manager is upgraded to 7.11.3.**

If you encounter the error, perform the following steps:

1. Deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.11.3.
2. Restart the cluster with a delay of 10 minutes.
3. Continue to upgrade CDH 7.1.8 to CDH 7.1.9.

**OPSAPS-68159: If you did not deactivate the Ozone parcel 718.2.x on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Ozone roles during the CDH 7.1.8 upgrade to CDH 7.1.9.**

If you encounter the error, perform the following steps:

1. Deactivate the Ozone parcel 718.2.x.
2. Restart the Ozone service.
3. Perform Finalize Upgrade for the Ozone service.

Step result: The Ozone roles will be displayed in green.

**CDPD-60989: The packaging version for Apache Ozone points to the 1.2.0 older version. This is a version string problem and not a packaging issue. The version of the Apache Ozone binary is closest to 1.4.0.**

None. This only affects the JAR versioning.

**CDPD-60366: Native library loader fails when system property `native.lib.tmp.dir` is not set. It fails because the library is copied to / instead of the cwd.**

Setting `native.lib.tmp.dir` to a certain path like `/tmp` should solve the issue. To set this system property, you must add `-Dnative.lib.tmp.dir=/tmp` to `ozone_java_opts` in Cloudera Manager configuration of the Ozone cluster.

**CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.**

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

**CDPD-60598: DataNode can consume up to the JVM heap size worth of direct memory buffers under specific failure scenarios. This can lead to a DataNode crash. This is due to Netty allocating and not freeing the direct memory buffers.**

Restarting DataNode will resolve this issue.

**CDPD-60466: The Ozone client is missing in the Cloudera Manager's Service Monitor node. This causes Ozone Canary's health check to fail.**

Ensure you install Ozone CLI on the same node as Cloudera Manager's Service Monitor.

**CDPD-60012: The Ozone SCM may be stuck in safe mode after upgrade to 7.1.9. This is due to a bug in how the SCM accounts for open storage containers.**

The workaround is to temporarily set the Ozone configuration `hdds.scm.safemode.threshold.pct` to a lower value like 0.90 and restart the SCM.

**CDPD-59679: In some scenarios, excess UNHEALTHY replicas of EC containers may not be removed from the cluster.**

None

**CDPD-50116: Topology-aware reads can provide better performance by routing applications to the closest replica of a file, if available, on the same physical node or same rack. However this feature is disabled by default in CDP 7.1.9.**

This feature can be enabled by setting `ozone.network.topology.aware.read` to true using Cloudera Manager safety valve.

**CDPD-57165: The DataNode disk usage thread may abort due to Ratis log tailing.**

Use `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory` to calculate disk usage. Configure in the `ozone-site.xml` property: `hdds.datanode.du.factory` value: `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory`

**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.

**CDPD-55513: Hive external table replication policy does not migrate Hive tables in HDFS storage to Ozone.**

To replicate data from HDFS to Ozone, see [Migrating your data from HDFS to Ozone](#). You can use HMS Mirror to replicate the metadata.

**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"`

**OPSAPS-67650: Ozone uses RocksDB as a library to persist metadata locally.**

By default, RocksDB places certain executables in `/tmp`, and thus encounters errors when `/tmp` is mounted with `noexec`.

The workaround is to configure RocksDB to put executables at another location. On a PhatCat node, the steps are:

1. Go to Cloudera Manager UI OZONE Configuration .

2. Find Ozone Service Environment Advanced Configuration Snippet (Safety Valve) and set the following environment variable: `ROCKSDB_SHARED_LIB_DIR=/VAR/TMP`
3. Restart Ozone.

**CDPD-49137: Ozone Manager Kerberos token expires for SCM communication and OM does not log in again.**

Sometimes, OM's Kerberos token is not updated and it stops to communicate with SCM. When this occurs, writes start failing.

Restart OM or set the safety valve `hadoop.kerberos.keytab.login.autorenewal.enabled = true`

**CDPD-56684: Keys get deleted when you do not have permission on volume**

When a volume is deleted, it recursively deletes the buckets and keys inside it and only then deletes the volume. The volume delete ACL check is done only in the end, due to which you may end up deleting all the data inside the volume without having delete permission on the volume.



**Note:** There was no bucket/key permission available which allowed the user to delete them recursively.

**CDPD-50610: Large file uploads are slow with OPEN and stream data approach**

Hue file browser uses the append operation for large files. This API is not supported by Ozone in 7.1.9 and therefore large file uploads can be slow or timeout from the browser.

Use native Ozone client to upload large files instead of the Hue file browser.

**OPSAPS-64097: Ozone service restart failed at SCM**

Stopping SCM service using Cloudera Manager can sometimes timeout and need a retry. The Cloudera Manager API waits for 90 seconds which is not sufficient under certain circumstances.

Retry the shutdown using Cloudera Manager if the SCM still shows as running after a refresh of the service status.

**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-filesystem-hadoop3.jar`, it is likely due to the presence of the user `hdfs` on the node prior to CDP parcel activation. As a result, the Cloudera Manager agent skips the initialization related to HDFS, leading to the non-creation of the `/var/lib/hadoop-hdfs` directory.

Create the directory `“/var/lib/hadoop-hdfs”` on all nodes followed by the deactivation and activation of the CDP parcel (deactivate and activate the Ozone parcel instead, in case Ozone parcel is used).

**CDPD-50447: When SCM High Availability is enabled, each of the SCM web UIs report the host of the web UI as the leader of HA, and the other two as followers. This gives wrong information**

Correct output is available by running the `ozone admin scm roles --service-id=<ID>` command.

**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](#).

**OPSAPS-62327: In an Ozone cluster without any gateway roles, Ozone is unable to deploy client configurations and displays the ConfigGenException error.**

You must add the Ozone gateway roles to the cluster.

**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-36389: The configurations `datanodes.involved.max.percentage.per.iteration` and `size.moved.max.per.iteration` are meant to limit the max number of DataNodes that is involved and maximum size that can move in an iteration. This bug causes balancer to stop an iteration when it is two DataNodes or 1 Container size (5GB) away from hitting these limits. However, these DataNodes can again be considered for balancing in the next iteration. This means the cluster is balanced after enough iterations, albeit a bit slowly. This bug is apparent in small clusters of around four DataNodes where the DataNodes can be either the source or target for a lot of moves but the iteration gets stopped when three DataNodes have been involved. It takes a higher number of iterations to eventually balance this cluster. While this is a performance issue, it does not prevent balancer from ultimately balancing the cluster.**

Increase the speed for balancing by decreasing the interval between each iteration using the configuration `balancing.iteration.interval`. The value of this configuration must be greater than `hdds.datanode.du.refresh.period`. `size.moved.max.per.iteration` can be increased to allow more data to move in one iteration.

**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-35141: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2> (state=08S01,code=40000) java.sql.SQLException: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2>. This issue does not occur if the source and target buckets are different in Hive queries. Currently, copying across the same bucket is only supported.**

Avoid different buckets in source and target path. Ensure that `hive.exec.stagingdir` is unset or set to the same bucket as the warehouse so that staging directories are created in the same bucket as the warehouse.

**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-41184: With Legacy buckets, FileSystem Optimised is not interoperating with the Ozone shell command. The cause is that the directory key entry in the DB KeyTable is stored as "dir1/" with trailing slash. But while performing the described operation, Ozone shell (o3://) is normalizing the given path and removing the trailing slash "/" from it. That results in KEY\_NOT\_FOUND exception.**

There are three workarounds:

- Use FileSystem API to Delete the Directories rather than Shell-Command API.
- Use FSO buckets instead of Legacy Buckets. As in FSO, you can create Intermediate Directories and Delete Directories using the Ozone shell commands.
- Disable and set the configuration `ozone.om.enable.filesystem.pathsflag` to false in order to delete the directories. This is generally not a preferred workaround because the cluster must be restarted to pick up the new changes.



**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-12542: Mount of Ozone filesystem with the help of FUSE fails.**

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-42897: EC writes are failing with "No enough datanodes to choose" after EC replication config is set globally.**

EC writes start failing when large number of pipelines are created as a result of multiple EC configs with different chunk sizes used to write keys.

If standard EC configs (that is, rs-3-2-1024k) are used to write keys, number of pipelines created per DataNode is limited to five and this issue is not seen with standard EC configs.

The recommendation is not to create too many random chunk sizes. It is configurable because, users can decide based on their workload. But do not to have separate chunk sizes for each file.

**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-40560: Filesystem Operations using Hadoop S3a connector on a `FILE_SYSTEM_OPTIMIZED` bucket fails. `org.apache.hadoop.ozone.om.exceptions.OMException: Unable to get file status: volume: s3v bucket: fso key: test/`**

Do not run Hadoop S3a commands on a FILE\_SYSTEM\_OPTIMIZED bucket. Use OBJECT\_STORE bucket layouts.

**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.

**CDPD-43432: Ozone Service in fault state in DataNode - Long Running setup.**

Upgrade RocksDB to the latest version.

**OPSAPS-63999: In the newly installed cluster, the Finish upgrade option is clickable.**

None

**OPSAPS-64648: Failed to start Ozone node using Cloudera Manager if default log path /var/log/hadoop-ozone does not exist. If this path does not exist, any Ozone nodes(for example SCM or DataNode) restart will fail.**

Run the following command `sudo -u hdfs mkdir -p /var/log/hadoop-ozone` or replace hdfs with the user Ozone roles that are running.

**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.

## Technical Service Bulletins

**TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager**

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.

## Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

## Known Issues in Ranger

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

## Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8

There are no new known issues identified in this release.

## Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

### UnsupportedClassVersionError

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.

### CDPD-68806: The Revoke operation for users belonging to a group or role permission does not function as expected.

List command is listing all the tables even when the user permission is revoked. And also the command does not add any deny policy to Ranger for that specific user.

For example:

user\_1 is the part of group\_1

user\_2 is the part of role\_1

Create a Ranger policy for group\_1 and role\_1 to access all the resources in HBase.

Try to revoke the permission for user\_1 and user\_2 using HBase shell global revoke command:

```
revoke 'user_1'
```

```
revoke 'user_2'
```

But there is no deny policy added for that user. And access is still granted to user\_1 and user\_2.

This behavior is currently not supported in HBase shell. Must be handled manually using the Ranger policy change.

### CDPD-68739: The revoke command does not work when using the HBase shell

While using the HBase shell, running the revoke command does not cancel the user permission. Users are able to perform actions even after running the revoke command.

For example:

When a user is granted permission to a particular table using the HBase grant command.

```
hbase:003:0> grant 'user_1', 'RWCA', 'some_abc_table'
```

And then execute the global revoke command for that user:

```
hbase:005:0> revoke 'user_1'
```

Permission for that user 'user\_1' for the table 'some\_abc\_table' not got revoked.

None.

### CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

### CDPD-58860: After upgrading CDP-7.1.8 to CDP-7.1.9, cdp-proxy token missing from knox-ranger policy.

As part of [OPSAPS-67480](#) in 719 default ranger policy is added from cdp-proxy-token topology, so that after a new installation of CDP-7.1.9, the knox-ranger policy includes cdp-proxy-token. However, upgrades do not add cdp-proxy-token to cm\_knox policies automatically.

Manually add cdp-proxy-token to the Knox policy, using Ranger Admin Web UI.

1. Log in to Cloudera Manager Ranger Admin Web UI, as a Ranger administrator.
2. On Ranger Admin Web UI Service Manager Resource Knox, click cm\_knox.
3. In Knox Policies, open the CDP Proxy UI, API and Token policy.
4. In Knox Topology\*, add cdp-proxy-token.
5. Click Save.
6. Restart Ranger.

**CDPD-58704: hadoop roll key / key delete command shows operation failed error when one KMS host is down, even when operation succeeds**

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered kms instances. if one KMS instance is stopped (not deleted), client gets 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.**

```
hbase:001:0> revoke 'hrt_11'
ERROR: org.apache.hadoop.hbase.coprocessor.CoprocessorException: HTTP 400 Error: processSecureRevokeRequest processing failed
 at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor.preRevoke(RangerAuthorizationCoprocesor.java:1309)
 at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor.preRevoke(RangerAuthorizationCoprocesor.java:1128)
 at org.apache.hadoop.hbase.master.MasterCoprocesorHost$162.call(MasterCoprocesorHost.java:1857)
 at org.apache.hadoop.hbase.master.MasterCoprocesorHost$162.call(MasterCoprocesorHost.java:1854)
 at org.apache.hadoop.hbase.coprocessor.CoprocessorHost$ObserverOperationWithoutResult.callObserver(CoprocessorHost.java:558)
 at org.apache.hadoop.hbase.coprocessor.CoprocessorHost.executeOperation(CoprocessorHost.java:631)
 at org.apache.hadoop.hbase.master.MasterCoprocesorHost.preRevoke(MasterCoprocesorHost.java:1854)
 at org.apache.hadoop.hbase.master.MasterRpcServices.revoke(MasterRpcServices.java:2740)
 at org.apache.hadoop.hbase.shaded.protobuf.generated.MasterProtos$MasterService$2.callBlockingMethod(MasterProtos.java)
 at org.apache.hadoop.hbase.ipc.RpcServer.call(RpcServer.java:387)
 at org.apache.hadoop.hbase.ipc.CallRunner.run(CallRunner.java:139)
 at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:369)
 at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:349)
```

None

**CDPD-56741: Improvement in log message when jwtauth not used**

The following exception is printed at startup only and it is not cluttering logs:

```
2023-05-30 06:18:40,127 ERROR org.apache.ranger.rms.security.RMSJwtAuthFilter:
```

```

quasar-pibgzl-1.quasar-pibgzl.root.hwx.site-startStop-1]:
Failed to initialize Ranger RMS JWT Auth Filter.
java.lang.Exception: RangerJwtAuthHandler:
Mandatory configs ('jwks.provider-url' & 'jwt.public-key')
are missing, must provide atleast one.
 at org.apache.ranger.authz.handler.jwt.RangerJwtAuthHan
dler.initialize(RangerJwtAuthHandler.java:84)
 ~[ranger-authn-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
 at org.apache.ranger.rms.security.RMSJwtAuthFilter.init
ialize(RMSJwtAuthFilter.java:77)
 ~[ranger-rms-common-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0
-186]
 at jdk.internal.reflect.NativeMethodAccessorImpl.invoke
0(Native Method) ~[?:?]
 at jdk.internal.reflect.NativeMethodAccessorImpl.invoke(N
ativeMethodAccessorImp

```

This exception is logged, since mandatory JWT auth filter configs are not provided (expected for some environments like y-cloud) Even though JWT auth filter is failed to initialised, it falls back to kerberos auth, hence no impact from auth perspective

NA

**CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup**

This is a warning log printed in catalina.out file when Ranger RMS server is initialised. Following exception is observed only in Oracle 19 setup. FileNotFoundException: /usr/share/java/oraclepki.jar

NA

**CDPD-48975: Ranger KMS KTS to KMS DB migration : keys with the same name but different case are not migrated**

KMS keys are not case sensitive

No workarounds. Such keys combination are very rare and migration doc was updated to check such keys before starting the migration.

**CDPD-42598: Kafka policy creation allowed with incorrect permissions.**

When creating a Kafka policy from the UI, the permissions "Idempotent write" and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource

None.

**CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload:**

```
{"resourceName": "classification", "userInput": "", "resources": {"classification": []}}
```

expectation is to return all the classifications . But the response has only "classification" Happens similarly for entity-label , entity-business-metadata.

None.

**CDPD-40734: User allowed to insert data into a hive table when there is a deny policy on a table column.**

A user is allowed to enter data into a table even if there is a deny policy present on one of the table columns.

Test scenario details:

Policy setup :-

```

policy 1 :- all access policy for hrt_qa, hive and impala
users
resources - database - * , table - *, column - *
users : hrt_qa, hive, impala
access - all access allowed
policy 2 :- policy on test_1.table_1 for hrt_5
users : hrt_5
resources : database - test_1, table - table_1, column - *
access :- all access allowed
policy 3 :- deny policy on test_1.table_1.c0 for hrt_5
users : hrt_5
resources : database - test_1, table - table_1, column -
c0
access - all access denied
data setup :-
database - test_1
table - table_1(c0 int, c1 int)

```

The user is able to insert data into the table.

None.

### Known Issues in Spark

Learn about the known issues in Spark the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8

The following known issues were identified in this release:

No new known issues in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

#### CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the RuntimeWarning: Failed to add file message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

#### CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

**CDPD-67517: Spark3 tests are failing if /tmp is mounted as NOEXEC.**

Map the tmpdir to a writable path in spark3-conf/spark-defaults.conf using the following steps:

1. Go to the Cloudera Manager.
2. Go to the Spark service.
3. Click the Configuration tab.
4. Select Scope Gateway .
5. Select Category Advanced .
6. Locate the Spark Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-defaults.conf\_client\_config\_safety\_valve property.
7. Map the tmpdir to a writable path:

```
spark.driver.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
spark.executor.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
```

8. Enter a Reason for change, and then click Save Changes to commit the changes.
9. Deploy the client configuration.

**CDPD-23817: In the upgraded Cluster, the permission of /tmp/spark is restricted due to the HDP configuration hive.exec.scratchdir=/tmp/spark.**

If you are using the /tmp/spark directory in the CDP cluster, you must provide the required additional Policies/ACL permissions.

**CDPD-22670 and CDPD-23103: There are two configurations in Spark, "Atlas dependency" and "spark\_lineage\_enabled", which are conflicted. The issue is when Atlas dependency is turned off but spark\_lineage\_enabled is turned on.**

Run Spark application, Spark will log some error message and cannot continue. That can be restored by correcting the configurations and restarting Spark component with distributing client configurations.

**CDPD-23007: Mismatch in the Spark Default DB Location. In HDP 3.1.5, hive\_db entities have one attribute - 'location' which is configured to the '/managed' path. In fresh install of CDP 7.1.7, hive\_db entities now have 2 attributes 'location' configured to '/external' path and 'managedLocation' configured to '/managed' path. In the AM2CM migration (HDP 3.1.5 -> CDP 7.1.7), the 'location' attribute from hive\_db entities in HDP 3.1.5 comes unaltered to CDP 7.1.7 and hence maps to '/managed' path.**

This issue arises only if you are upgrading from HDP 3.1.5 to CDP 7.1.7. If you are performing a fresh install of CDP 7.1.7, you can ignore this issue.

None

**CDPD-217: The Apache Spark connector is not supported**

The old *Apache Spark - Apache HBase Connector* (shc) is not supported in CDP releases.

Use the new HBase-Spark connector shipped in CDP release.

**CDPD-3038: Launching pyspark displays several HiveConf warning messages**

When pyspark starts, several Hive configuration warning messages are displayed, similar to the following:

```
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of name hive.meta
store.runworker.in does not exist
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of
name hive.masking.algo does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of
name hive.metastore.runworker.in does not exist
```

```
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of name hive.masking.algo does not exist
```

These errors can be safely ignored.

### Known Issues in Sqoop

Learn about the known issues in Sqoop, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8

#### CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version

If you have upgraded to CDP PvC Base 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to CDP PvC Base 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see [Troubleshooting Apache Sqoop issues](#).

#### CDPD-44431: Using direct mode causes problems

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the `--direct` option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the `sqoop.enable.deprecated.direct` property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through `-Dsqoop.enable.deprecated.direct=true`.

#### CDPD-3089: Avro, S3, and HCat do not work together properly

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

#### Parquet columns inadvertently renamed

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.

#### PARQUET-99: Importing Parquet files might cause out-of-memory (OOM) errors

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.

### Known Issues in Schema Registry

Learn about the known issues in Schema Registry, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8

There are no new known issues identified in this release.



**Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8****CDPD-40380: Authorization checking issue when Kerberos is disabled**

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the `ranger.plugin.schema-registry.service.name` field.
3. Replace `GENERATED_RANGER_SERVICE_NAME` with the actual name of the service.
4. Restart the Schema Registry service.

**CDPD-49304: AvroConverter does not support composite default values**

AvroConverter cannot handle schemas containing a `STRUCT` type default value.

None.

**CDPD-59015: Schema Registry does not create new versions of schemas even if the schema is changed**

Schema Registry uses a schema fingerprinting mechanism to differentiate between schemas. However, fingerprinting does not take into consideration the schema attributes of the field type. As a result, if you have two schemas where the only difference is that one has type attributes defined and the other does not, they will be considered identical by Schema Registry. For example, the following schemas are considered identical:

```
#Schema V1
{"type":"record","name":"schema_name","namespace":"ns","fields":
[{"name":"local_timestamp_micros_long","type":"long"}]}

#Schema V2
{"type":"record","name":"schema_name","namespace":"ns","fields":
[{"name":"local_timestamp_micros_long","type":{"type":"long",
"logicalType":"local-timestamp-micros"}]}]
```

Notice that the only difference is that in the second schema, the `local_timestamp_micros_long` field has a logical type specified. In cases like this, the new version of the schema is not created, the initial version is used. This is true even if the data that is being produced has a new schema version. The ID of the first schema version is used and is put in the serialized record. The new schema version is not created.

This issue is common when using change data capture (CDC) connectors like the Debezium connectors. This is because CDC connectors create schemas with the logical type decimal based on the column type in the database schema. For example:

```
{"type":"record","name":"schema_name","namespace":"ns","fields":
[{"name":"database_column","type":{"type":"bytes","logicalType":
"decimal","precision":64,"scale":0}}]}
```

If the database schema changes (for example, the column type), it is possible that only scale changes, which is a schema attribute.

```
{"type":"record","name":"schema_name","namespace":"ns","fields":
[{"name":"database_column","type":{"type":"bytes","logicalType":
"decimal","precision":64,"scale":1}}]}
```

In this case, even though scale changed to 1, the first version of the schema is used where scale is 0. As a result, the data is consumed with the wrong scale.

Avoid using logical types or other attributes. Alternatively, ensure that there are no changes in the logical types or other attributes between schema versions.

**OPSAPS-68708: Schema Registry might fail to start if a load balancer address is specified in Ranger**

Schema Registry does not start if the address specified in the Load Balancer Address Ranger property does not end with a trailing slash (/).

Set the value of the RANGER\_REST\_URL Schema Registry environment variable to an address that includes a trailing slash.

1. In Cloudera Manager, select the Schema Registry service.
2. Go to Configuration.
3. Find the Schema Registry Server Environment Advanced Configuration Snippet (Safety Valve) property and add the following:

```
Key: RANGER_REST_URL
Value: [***RANGER REST API URL***]
```

Replace [\*\*\*RANGER REST API URL\*\*\*] with an address that can be used by Schema Registry to access Ranger. Ensure that the address ends with a trailing slash. For example: `https://ranger-1.cloudera.com:6182/`

4. Restart the Schema Registry service.

#### **OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade**

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

1. Access the Ranger Console (Ranger Admin web UI).
2. Click the `cm_atlas` resource-based service.
3. Add the `schemaregistry` user to the `all - *` policies.
4. Click `Manage Service Edit Service`.
5. Add the `schemaregistry` user to the `default.policy.users` property.

#### **OPSAPS-69317: Kafka Connect Rolling Restart Check fails if SSL Client authentication is required**

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to `required`. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to `requested` instead of `required` and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

#### **Known Issues in Streams Messaging Manager**

Learn about the known issues in Streams Messaging Manager, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8**

There are no new known issues identified in this release.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8**

##### **OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager**

Cloudera Manager does not display a Log Files menu for SMM UI role (and SMM UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by SMM UI is not supported by Cloudera Manager.

View the SMM UI logs on the host.

##### **OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners**

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL\_SSL). You would have to override bootstrap server URL (hostname:port as set

in the listeners for broker). Add the bootstrap server details in SMM safety valve in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Add the following value for bootstrap servers>Save Changes > Restart SMM.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

#### **CDPD-70313: KNOX does not send Authentication header on FIPS configuration**

KNOX does not send neither the Authentication header nor the `hadoop.auth` cookie. Because of this, the SMM UI returns an HTTP 401 response and sets the `www-authenticate: Negotiate` header. After this, KNOX still does not send the Authentication header. This results in the SMM UI being inaccessible through Knox.

You can access the SMM UI directly in Cloudera Manager at `Clusters SMM Streams Messaging Manager Web UI` and log in using the Kerberos username and password.

#### **CDPD-82560: Streams Messaging Manager known unprotected endpoint**

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**

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**

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.

#### **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 in Cloudera Runtime 7.1.9 SP1 CHF 8.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8**

There are no new known issues identified in this release.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8**

##### **CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic**

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

##### **CDPD-11079: Blacklisted topics appear in the list of replicated topics**

If a topic was originally replicated but was later excluded for replication, 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, this topic will be included in the response. Additionally, the excluded topic will also be visible in the SMM UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

#### **CDPD-30275: SRM may automatically re-create deleted topics on target clusters**

If `auto.create.topics.enable` is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with `srm-control`. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1]
```

2. Wait until SRM is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

#### **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 in Cloudera Runtime 7.1.9 SP1 CHF 8.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8**

There are no new known issues identified in this release.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8**

##### **COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."**

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

None.

##### **COMPX-4644: Queue capacity rounding problem when configuration is initially set via YARN**

When setting the capacity scheduler configuration through the YARN/Cloudera Manager configuration, there may be capacity values that use multiple decimal places. This results in rounding/floating point precision discrepancies in the UI when trying to validate that all sibling capacities equal 100%. The UI looks like all the numbers add up to 100, but the validation still displays an error and does not allow to save the capacities. It is also observed that the capacity is being calculated as, for example, 99.9999999991 in the backend.

- Create queues within the UI, or
- Ensure that capacities configured through the Capacity Scheduler safety valve do not have more than one decimal place.

#### **20202 Database migration after enabling opt-in migration**

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to CDP 7.1.9, forcing migration from H2 to PostgreSQL database.
- Upgrade to CDP 7.1.9 CHF2, moving back to H2 database.
- Upgrade to CDP 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version CDP 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

#### **CDPD-56559: MapReduce jobs can intermittently fail during a rolling upgrade.**

During a rolling upgrade between CDP versions 7.1.8 and 7.1.9, MapReduce jobs may fail with message, `RuntimeException: native snappy library not available`. Although the native Snappy compression library is not loaded, a checkmark displays to indicate that the Snappy compression library is loading for NodeManagers that are pending upgrades. This causes the MapReduce jobs that are associated with the NodeManagers to fail. After the upgrade, the jobs work as expected. This issue only impacts rolling upgrades from before CDP 7.1.9 to a higher version.

None.

#### **COMPX-12021 Queue Manager configurations on Scheduler Configuration page are not working**

When setting the following properties on the YARN Queue Manager UI, the properties are set in the `capacity-scheduler.xml` file which does not have any effect on YARN. The properties need to be set in the `yarn-site.xml` file, which does not happen when you set them through YARN Queue Manager.

- "Maximum Application Priority" – `"yarn.cluster.max-application-priority"`
- "Enable Monitoring Policies" – `"yarn.resourcemanager.scheduler.monitor.enable"`
- "Monitoring Policies" – `"yarn.resourcemanager.scheduler.monitor.policies"`



**Note:** You can also set this property on the YARN configuration page in Cloudera Manager as "Capacity Scheduler Preemption".

- "Preemption: Observe Only" – `"yarn.resourcemanager.monitor.capacity.preemption.observe_only"`
  - "Preemption: Monitoring Interval (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.monitoring_interval"`
  - "Preemption: Maximum Wait Before Kill (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.max_wait_before_kill"`
  - "Preemption: Total Resources Per Round" – `"yarn.resourcemanager.monitor.capacity.preemption.total_preemption_per_round"`
  - "Preemption: Over Capacity Tolerance" – `"yarn.resourcemanager.monitor.capacity.preemption.max_ignored_over_capacity"`
  - "Preemption: Maximum Termination Factor" – `"yarn.resourcemanager.monitor.capacity.preemption.natural_termination_factor"`
  - "Enable Intra Queue Preemption" – `"yarn.resourcemanager.monitor.capacity.preemption.intra-queue-preemption.enabled"`
1. In Cloudera Manager, select the YARN service.
  2. Click the Configuration tab.
  3. Search for `yarn-site.xml`.
  4. Under YARN Service Advanced Configuration Snippet (Safety Valve) for `yarn-site.xml`, add the corresponding parameter and value you need.
  5. Click Save Changes.
  6. Restart the YARN services.

#### **COMPX-13177: QueueManager webapp requests fail with 'HTTP ERROR 400**

**`java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA'`**

Products:

- Cloudera Manager for CDP Private
- Cloud Base Cloudera Manager for CDP Public Cloud

Context:

- Centos 7.8 and Redhat 7.8 operating systems, when FIPS support is enabled.

Problem:

- When attempting to display the Yarn Queue Manager interface, Cloudera Manager displays an error: "HTTP ERROR 400 java.net.ConnectException: Unsupported cipher suite 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.

### **Third-party applications do not launch if MapReduce framework path is not included in the client configuration**

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the `mapreduce.application.framework.path` property is set to the appropriate value, but third-party applications with their own configurations does not launch.

Set the `mapreduce.application.framework.path` property to the appropriate configuration for third-party applications.

### **JobHistory URL mismatch after server relocation**

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated the `mapred-site.xml` file that references the correct JobHistory Server.

### **YARN cannot start if Kerberos principal name is changed**

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN does not start. In such cases, the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the `znode` and restart the YARN service.
- Use the `reset ZK ACLs` command. This also sets the `znodes` below `/rmstore/ZKRMStateRoot` to `world:anyone:cdw` which is less secure.

### **Queue Manager does not open on using a custom user with a default Kerberos principal**

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager `process_username` property matches the YARN `process_username` property.

### **Known Issues in Zeppelin**

Learn about the known issues in Zeppelin the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8**

There are no new known issues identified in this release.

### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8**

**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.

#### **BUG-125263: Zeppelin service move fails on clusters upgraded from HDP3.1.5**

Resolve the circular symlink issue on the Zeppelin node by linking the conf directory to a new directory under /etc/zeppelin:

- # mkdir -p /etc/zeppelin/<version>/0
- # rm /usr/hdp/<version>/zeppelin/conf
- # ln -s /etc/zeppelin/<version>/0 /usr/hdp/<version>/zeppelin/conf

Where version is the HDP version. For example, 7.1.x-yyy. Restart the Zeppelin server in Ambari.

#### **CDPD-3090: Due to a configuration typo, functionality involving notebook repositories does not work**

Due to a missing closing brace, access to the notebook repositories API is blocked by default.

From the CDP Management Console, go to Cloudera Manager for the cluster running Zeppelin. On the Zeppelin configuration page (Zeppelin serviceConfiguration), enter shiro urls in the Search field, and then add the missing closing brace to the notebook-repositories URL, as follows:

```
/api/notebook-repositories/** = authc, roles[{{zeppelin_admin_group}}]
```

Click Save Changes, and restart the Zeppelin service.

#### **CDPD-2406: Logout button does not work**

Clicking the Logout button in the Zeppelin UI logs you out, but then immediately logs you back in using SSO.

Close the browser.

#### **OPSAPS-59802: Zeppelin and Livy roles should be co-located on the same host.**

When installing or upgrading to CDP Private Cloud Base, you must co-locate all Zeppelin and Livy roles on the same cluster host due to an issue with certificate generation.

#### **Known Issues in ZooKeeper**

Learn about the known issues in ZooKeeper the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 8.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 8**

There are no new known issues identified in this release.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 8**

There are no new known issues identified in the previous releases.

### **Cumulative hotfix 7**

Know more about the cumulative hotfix 7 for CDP 7.1.9 SP1. This cumulative hotfix was released on 07 Apr, 2025.

#### **Cloudera Runtime 7.1.9.1037 (Cumulative Hotfix 7) download URL:**

##### **Parcel Repository Location**

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1037/parcels/
```

#### **Fixed issues in 7.1.9 SP1 CHF 7**

Know more about the cumulative hotfix 7 for 7.1.9 SP1.



**Note:** Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p1037.64780234.

**CDPD-78671: Metric timer task is blocking installSnapshotFromLeader on follower node**

ozone.om.snapshot.rocksdb.metrics.enabled is now available in Ozone to disable metric collection on the snapshotted DB if necessary.

**CDPD-73375: Publishing hadoop metrics immediately in Prometheus sink fills up SinkQueue quickly**

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.

**CDPD-80823: Snapshot creation is removing extra keys from Active Object storage's DB**

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.

This issue is fixed and the extra keys will not be removed from the DeletedTable for Active Object storage.

**CDPD-78781: Tarball creation is interfering with snapshot purge**

If Tarball creation is in process while the snapshot is getting purged, it fails the snapshot db dir delete command causing the snapshot db dir to linger around even though the snapshot is purged from the snapshotInfoTable and needs to manually delete the dir.

This issue is fixed and the Synchronized SnapshotDeletingService on BootstrapStateHandler.Lock makes sure that no background service is running when Tarball creation is in progress.

**CDPD-57559: New Ozone Manager leader cannot verify the Ozone delegation token signed by old Ozone Manager leader**

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.

**CDPD-81557: Improve Ranger Admin Diagnostic Collection command from Cloudera Manager scripts**

The Ranger Admin Diagnostic Collection command is now enhanced and a new configuration option called `ranger.admin.diag.metrics.collection.type` is now introduced. This option allows you to specify the type of metrics data to be collected. This option appears while you are upgrading from 7.1.9 SP1 CHF6 to 7.1.9 SP1 CHF7, and upgrade Cloudera Manager.

**CDPD-80112: Knox service might fail due to JARs picked up from the /usr/share/java folder**

Knox service might fail due to Java Archive (JAR) files picked up from the /usr/share/java folder.

This issue is now resolved.

**Apache JIRA:** [KNOX-3108](#)

**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.

**CDPD-80628: Unable to query Iceberg tables from Impala**



After upgrading to CDP Private Cloud Base version 7.1.9 SP1 CHF 6 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](#)

#### **CDPD-79911: Netty upgrade to 4.1.118.Final**

Upgraded Netty to version 4.1.118.Final due to CVE-2025-24970, CVE-2025-25193.

#### **CDPD-77107: 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](#)

#### **CDPD-80208: When the ignore pattern property is added on a hook, Iceberg tables are also ignored**

When using the `atlas.hook.hive.hive_table.ignore.pattern=.*test_.*` property as an ignore pattern by Hive Server 2 or Hive Metastore hooks, it no longer ignores `iceberg_table` entities next to Hive entities. `hive_table` ignore patterns ignore only `hive_table` related entities.

#### **CDPD-72496: File extension restrictions for Hue file uploads**

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-77746: Metrics Collection Failure with "DPI-1010: not connected" on Oracle DB backend**

Metrics collection fails with the error "DPI-1010: not connected" when using an Oracle database as the backend.

This issue is now fixed.

#### **CDPD-79237: Hive Metastore schema upgrade fails due to NULL values**

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.

**CDPD-81079: Metastore did not enforce maximum Thrift message size**

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 is addressed by applying the configured maximum Thrift message size on the server, ensuring consistent behavior and avoiding unexpected disconnections.

Apache Jira: [HIVE-28824](#)

**CDPD-80942: Aborted transactions due to timeout are not logged in notification log**

Transactions aborted due to timeout are not logged as ABORTED in the NOTIFICATION\_LOG.

To address the issue, the fix includes logging an ABORT event for transactions aborted due to timeout. This allows such transactions to be replicated and cleaned up from the target cluster without delay, reducing unnecessary overhead and improving query performance.

Apache Jira: [HIVE-27797](#)

**CDPD-79990: Incorrect error shown for successful LOAD DATA command in Beeline**

When running a LOAD DATA statement using EXECUTE IMMEDIATE in Beeline with HiveServer2, the data loads successfully into the target table, but an error is incorrectly displayed in the console. HiveServer2 logs show a NullPointerException when generating result set metadata, even though the command completes and the table is correctly updated.

This issue is fixed.

Apache Jira: [HIVE-28766](#)

**CDPD-77626: Improving performance of ALTER PARTITION operations using direct SQL**

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](#)

**CDPD-49323: INSERT statement does not respect Ranger policies for HDFS**

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm\_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in

```
AnalysisException: Unable to INSERT into target table (default.t
1) because Impala does not have WRITE access to HDFS location: h
dfs://XXXXXXXXXXXXX
```

To address the issue, the fix includes skipping the HDFS permission check and assuming the Impala service user has read and write access to all HDFS paths associated with the target table during query analysis when Ranger is enabled. This ensures that Ranger policies are respected during insert operations.

Apache Jira: [IMPALA-11871](#)

**COMPX-77765: Hadoop - Upgraded Kafka-clients to version 3.7.2/3.8.1+ due to CVE-2024-56128**

Upgraded the Kafka client to version 3.9.0 to fix CVE-2024-31141.

**COMPX-78928: Backport HADOOP-16299 to avoid module access violation in LdapGroupsMapping**

When LdapGroupsMapping was the selected group mapping method, running any commands that relied 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 now avoided in the group mapping implementation without functional changes.

Apache Jira: [HADOOP-16299](#)

#### **CDPD-79763: Fix clobbering of files across epochs in Spark Structured streaming with Iceberg**

A bug in structured streaming was found that resulted in clobbering of files in Iceberg tables.

The issue was addressed by backporting an upstream fix.

#### **CDPD-77461: Livy: Upgrade mina-core to 2.0.27/2.1.10/2.2.4+ due to CVE-2024-52046**

Upgrades mina-core to patch a vulnerability issue.

#### **CDPD-78656: Health test for Knox fails if the gateway.client.auth.needed = true is set**

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
```

Also, set gateway.port.mapping.health to a value different than the gateway\_port to map the health test to another port.

For more information on excluding the topology, see the [Apache Knox Documentation](#).

There are no Common Vulnerabilities and Exposures (CVEs) fixed in this CHF.

### **Behavioral changes in Cloudera Runtime 7.1.9 SP1 CHF 7**

Know more about the behavioral changes in cumulative hotfix 7 for 7.1.9 SP1.

#### **Impala**

##### **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](#)

### Known issues in 7.1.9 SP1 CHF 7

The following section lists the known issues and limitations, the areas of impact, and workaround for each Runtime component.

#### Known Issues in Atlas

Learn about the known issues in Ozone the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7

The following known issues were identified in this release:

#### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7

##### CDPD-77054: Exporting shell entities can cause the import to fail

For Apache Hive tables created with Apache Spark, there may be shell entities created in Apache Atlas. If these entities are in the export zip file before they are resolved, the import will fail.

Workaround: None

##### CDPD-70450: Impala SQL queries that include the “WITH” clause should populate lineage in Atlas

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-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

##### CDPD-6565: Issue with ddlQueries in Atlas for Hive/Impala tables created in default/non-default db using DAS/HUE respectively

ddlQueries are not created for Impala origin table. ddlQueries created for Impala/Hive CTAS tables have difference in names as compared to same operation for Hive table. The Impala table has the query in the ddlQueries entity name.

None

##### CDPD-55301: ddlQueries and ALTERNATE\_\* lineage missing for Spark tables created through spark3-shell

The ddlQueries and ALTERNATE\_\* lineage missing for Spark tables created using spark3-shell.

None

##### CDPD-56085: 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

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*`.

None

**CDPD-58554: Discard audits of specific classification, label, and business metadata**

Support to control audits for specific classification, label, and business metadata is not present in the Custom Audit Filters feature.

None

**CDPD-40346: The `ddlQueries` and `ALTERTABLE_ADDCOLS` lineage missing for Impala tables.**

The `ALTERTABLE_ADDCOLS` lineage has some issue when an Impala table is altered and the corresponding lineage is not created.

**CDPD-67112: Import transforms do not work as expected when replacing a string which already has ":"**

The character ":" is not supported in path replacements. The import succeeds but location remains unchanged. The character ":" must be avoided.

None

**CDPD-69150: Unable to add labels or user defined properties in Japanese**

Adding Japanese labels or user defined properties results in the error message: "Invalid label: ###, label should contain alphanumeric characters, \_ or -"

None

**CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties**

Entities, such as `hdfs_path`, are not returned by the search when their indexable properties are searched using partial search terms made of Japanese or Chinese characters. Only exact matches return results when searching the indexed properties made of Japanese or Chinese characters.

None

**CDPD-68191: Suggestions do not return the correct results when searching multiple Chinese characters**

Free text search does not return results when searching Chinese phrases made of multiple characters. Partial searches return the correct results.

None.

**CDPD-71219: Regression : Suggestions don't work for single character words on indexable attributes**

When searching for entities whose name (entity names are indexable) is a single character, search results are returned but suggestions are not. When searching for entities whose description (entity descriptions are not indexable) is a single character, both search results and suggestions are returned.

None.

**CDPD-67450: Table name renaming operation is not updating or creating `iceberg_table` entity**

Renaming an Iceberg table does not update the corresponding Atlas entity.

None.

**CDPD-67089: Export/Import: When a table with Ozone path is exported as "connected", only the Ozone key is exported.**

When table with Ozone path is exported as "connected", only Ozone key is exported. Other Ozone entities, such as Ozone volume, Ozone bucket are not exported.

None.

**CDPD-43772: Performance issues with Atlas service**

If there are lot of update operations and the compression type of column families of atlas\_janus table is SNAPPY, then the Kafka message processing might become slower.

- Consider setting compression type of column families of atlas\_janus table as GZ.

**OPSAPS-67783: During rolling upgrade one among two Atlas server failed to start but Cloudera Manager considered as success**

Cloudera Manager marks the Execute command Start on service Atlas-1 as a success even when Atlas service had failed to start successfully. In such cases, Atlas logs give the exact reason for Atlas start-up failure.

None

**CDPD-19358: "IsIndexable" and "isOptional" value of a typedef's attribute is modified post migration.**

In HDP-265, falcon\_feed\_creation has "stored-in" attribute which has IsIndexable value set to True.

Post migration, "stored-in" is moved to relationshipAttributeDefs and has IsIndexable value set to False.

Similarly, in HDP-265 the table attribute of hive\_storage\_desc isOptional is set to true. Post migration, isOptional set to False.

None.

**CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command**

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**

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**

In converting content from Navigator to Atlas, the create time for Hive tables is not moved to Atlas.

None

**CDPD-10574: Suggestion order doesn't match search weights**

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**

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@58112
bc4 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) does not 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 the metadata collection from HMS. For example, if an ALTER TABLE operation in Spark that is changing a table name and is reporting 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 "@" does not 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, and use 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 does 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 logout 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 CDP.

None.

**CDPD-1664: Guest users are redirected incorrectly**

Authenticated users logging in to Atlas are redirected to the CDP 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: The 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, to add metadata tags to relationships that you wanted to make sure were unique in the system, you can design the relationship attribute with the property IsUnique equal true. However, in this release, the IsUnique attribute is not enforced.

None.

**DOCS-13759: Tag Propagation stops after a certain depth while the lineage is being extended**

When a tag is added to an entity at timestamp T1, the entities along the lineage to which the tag must be propagated is calculated at T1. Before tag propagation completes, if the lineage is extended, tag does not propagate to the entities in the extended lineage.

**CDPD-41142: When a Kafka console consumer group is run, more than one update audits are seen**

After running the console consumer with a consumer group, verify the consumer group entity created, along with the metrics and notifications for the consumer group and topic. The expected result can be: one ENTITY\_CREATE audit and one ENTITY\_UPDATE audit. But more than one ENTITY\_UPDATE audits are seen.

**CDPD-40165: Two audits are created for SPARK CTAS table**

When following Spark queries are fired:

```
spark.sql("create table table1(id int)")
spark.sql("create table table2 as select * from table1")
```

HMS sends "ENTITY\_CREATE" and "ENTITY\_FULL\_UPDATE\_V2".

The extra ENTITY\_FULL\_UPDATE\_V2 message received from HMS is sent as a part of ALTER\_TABLE\_ADDCOLS event from the HMS Hook side. This behaviour is observed only when the queries are run from Spark SQL and not when run the same queries are run from Beeline.

**CDPD-39197: Debug metrics returns empty data**

When debug metrics is enabled and some operations are performed, the response is empty

**CDPD-36495: Updating legacyAttribute from False to True resets the initially created relationshipAttributes values**

Creating types, entities, and to start, you must set the relationship with is\_legacy\_attribute value as False.

Later, update the value relationshipDef is\_legacy\_attribute to True.

For the entities that were created before updating the is\_legacy\_attribute to True, the relationship Attributes value is reset.

**CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities**

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

**CDPD-22744: Bulk entity DELETE API does not delete authorised entities**

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

**CDPD-29409: Hive import: Suggestion suggests entity which is deleted.**

Suggestions suggests tables of a database, which is a deleted entity.

**CDPD-25152: Tag propagation through deferred actions consumes additional time as compared to default flow**

The additional time might be due to the small overhead added to create / update task vertex and which is run in the background. This also depends on number of tasks queued to be executed.



**CDPD-42954: Zeppelin notebook fails after enabling Atlas-HDFS hook**

The Zeppelin notebooks are failing with errors after enabling Atlas-HDFS hook in the CDP cluster.

When the below properties are set for atlas-client.properties in Cloudera Manager:

- atlas.jaas.KafkaClient.option.keyTab
- atlas.jaas.KafkaClient.option.principal

Along with adding the properties in /etc/atlas/conf/atlas-application.properties, Cloudera Manager also adds these properties to atlas-application.properties for other services (like Spark).

Adding these properties interferes with the normal flow of the services (like Spark)

To enable HDFS lineage feature, instead of setting these properties through Cloudera Manager, users can manually add the properties directly in /etc/atlas/conf/atlas-application.properties

**CDPD-10576: Deleted Business Metadata attributes appear in Search Suggestions**

Atlas search suggestions continue to show Business Metadata attributes even if the attributes have been deleted.

None.

**CDPD-39427: [HDFS Lineage]: When the input is a directory in case of put/copyfromLocal/cp/mv commands, lineage is not created even though the script succeeds.**

When Source is a directory and target is a directory which is already present in Atlas, the command succeeds and inserts the data in the desired location, but lineage is not created.

**DOCS-13760: System Attributes search, \_\_classificationNames: Search with parent tag does not return entities associated to its children tags**

System attribute search with \_\_classificationNames = parent\_tag returns entities associated to parent\_tag only and not entities associated to its children tag.

Workaround: Instead of using system attribute, employ the basic search attribute "classification" which lists entities associated with inherited classifications.

**CDPD-35818: Basic search with tag filter provides approximateCount as -1 when there is no match and is 0 otherwise**

When the following search operations are performed:

- Faceted search with both tag filter and entity filter

The observed approximateCount is -1.

Here both entity filter and tag filter are present and when there is no match the response received is -1.

- Faceted search with only entity filter

Performing a basic search provides an approximate value of 0 when there is no match.

- Faceted search with only tag filter

Whenever there is tag filter in the query and there is no entity match, the approximateCount is -1 and if the tag filter is not available, the response approximateCount is 0

**CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities**

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

**CDPD-22744: Bulk entity DELETE API does not delete authorised entities**

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

**CDPD-76035: Resource lookup for Atlas service is failing**

Once the Atlas configuration snippet `atlas.authentication.method.file` is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

#### **CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout**

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

#### **CDPD-79160: NPE while deleting BusinessMetadata**

If business metadata is created without adding any applicable types, a `NullPointerException` is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

#### **CDPD-77911: Missing Log4j Redactor dependency**

The missing `org.cloudera.logredactor` dependency can result impossible Log4j errors during startup. Additionally, it can lead to unwanted data leak because of incorrect log redaction.

#### **CDPD-80921: Without permission for one glossary, /glossary call throws exception**

When you don't have the permission for a glossary, the `/glossary` call results in a 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-82056: UI: when server response date fields as '0', UI shows as current time**

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](#)

#### **CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export**

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**

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](#)

#### **CDPD-66938: [Analyze] [Atlas] test\_time\_range tests fail**

When the Apache Atlas server is running on a node which has time zone other than UTC, there might be a time of day when the search results might differ if the relative CreateTime date range filters of TODAY, YESTERDAY, etc. are used. 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.

Use explicit date range filters instead of using relative date range filters, such as, TODAY, YESTERDAY.

#### **CDPD-84502: Advanced Search is not working properly in Atlas UI**

The Advance Search in Atlas does not return any results when selecting the Typedef option. The API is not triggered for each relation.

None

**CDPD-76789: Creating tag with name description throws java.lang.ClassCastException**

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails**

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

**CDPD-70321: Atlas Parallel import is failing with various errors**

During a parallel import-export activity with six iceberg table policies with exportOption as db1.\* for all six exports, all import fail after the exports.

**CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen**

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

**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**

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn**

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

**CDPD-79099: Unable to import hive tables using import script in Public Cloud**

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-88230: Bulk import operations currently do not construct lineage as expected**

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**

Bulk import operations do not create audits for the imported entities.

**Known Issues in Cruise Control**

Learn about the known issues in Cruise Control, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

**Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7**

There are no new known issues identified in this release.

**Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7****Rebalancing with Cruise Control does not work due to the metric reporter failing to report the CPU usage metric**

On the Kafka broker, the Cruise control metric reporter plugin may fail to report the CPU usage metric.

If the CPU usage metric is not reported, the numValidWindows in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.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 the [Moving a Host Between Clusters](#).



**Note:** Cluster nodes affected by this issue are not displayed as unhealthy.

**Known Issues in HBase**

Learn about the known issues in HBase, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

**Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7****CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress**

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL

- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

#### **IntegrationTestReplication fails if replication does not finish before the verify phase begins**

During IntegrationTestReplication, if the verify phase starts before the replication phase finishes, the test will fail because the target cluster does not contain all of the data. If the HBase services in the target cluster does not have enough memory, long garbage-collection pauses might occur.

Workaround: Use the -t flag to set the timeout value before starting verification.

#### **HDFS encryption with HBase**

Cloudera has tested the performance impact of using HDFS encryption with HBase. The overall overhead of HDFS encryption on HBase performance is in the range of 3 to 4% for both read and update workloads. Scan performance has not been thoroughly tested.

Workaround: N/A

#### **AccessController postOperation problems in asynchronous operations**

When security and Access Control are enabled, the following problems occur:

- If a Delete Table fails for a reason other than missing permissions, the access rights are removed but the table may still exist and may be used again.
- If `hbaseAdmin.modifyTable()` is used to delete column families, the rights are not removed from the Access Control List (ACL) table. The `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.

Workaround: N/A

Apache Issue: [HBASE-6992](#)

#### **Snappy compression with /tmp directory mounted with noexec option**

Using the HBase client applications such as `hbase hfile` on the cluster with Snappy compression could result in `UnsatisfiedLinkError`.

Workaround: 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.

#### **The custom JWT Principal claim is not usable for HBase**

The `hbase.security.oauth.jwt.token.principal.claim` configuration property does not exist in the 7.1.9 SP1 version. The purpose of this property is to allow the use of a Subject or Principal claim different from the default sub claim.

Workaround: Use the standard sub claim in the JWT token.

#### **HBase procedures might hang because WriteAheadLog for ProcedureStore does not progress**

HBase Master might hang and cause longer recovery when the WriteAheadLog is stuck while persisting to the ProcedureStore. HBase Master eventually aborts itself, however, the operations that rely on the ProcedureStore (for example, region movements and table creation) do not progress during this period.

Workaround: Restart the active HBase Master to speed up the recovery.

**CDPD-77399: HBase fails to register the servlet metrics and throws `ClassNotFoundException: org.apache.hadoop.metrics.MetricsServlet`**

The `MetricsServlet` class is a Hadoop 2-based metric servlet unavailable in Hadoop 3 deployments.

Workaround: Ignore this WARN log message during HBase Master and RegionServer startup.

Apache Issue: [HBASE-28315](#)

**Known Issues in Hive**

Learn about the known issues in Hive, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

**Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7**

There are no new known issues identified in this release.

**Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7**

**CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress**

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

**CDPD-60770: Beeline Authentication Issue with Special Characters in Passwords**

When LDAP is enabled, users cannot authenticate with Beeline if the password contains a special character. For example, the following string fails:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>;user=user@XXX;password='R3G#xpXyoy1M0Jb1'
```

Use the `-p` parameter to execute the Beeline command:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>; -n user@XXX -p 'R3G#xpXyoy1M0Jb1'
```

**CDPD-43769: Inconsistent behavior with CHAR comparisons and string literals**

In Hive, 'char' data to 'string' data comparison is considered as equal, but 'char' data to 'varchar' data comparison is considered as not equal if the data has trailing spaces. For example,

```
create table test(colchar char(10), colvarchar varchar(20), colstring string);
insert into test values ('a ', 'a ', 'a ');
```

```
select count(*) from test where colchar = colstring;
Output returns 1 because char and string data is considered as
equal

select count(*) from test where colchar = colvarchar;
Output returns 0 because char and varchar data is considered as
not equal
```

It is recommended that you use 'varchar or string' data type instead of 'char' data type for any character or string representation.

#### **CDPD-42726: Wrong results for agg queries from stats after running multi-insert query on managed table**

Multiple inserts at the same time into the same partition results in invalid stats. For example:

```
create table source(p int, key int, value string);
insert into source(p, key, value) values (101,42,'string42');

create table stats_part(key int,value string) partitioned by (p
int);
from source
insert into stats_part select key, value, p
insert into stats_part select key, value, p;

select count(*) from stats_part;
-- In this scenario, StatsOptimizer helps serving this query be
cause the result should be rowNum of the partition p=101.
-- The result is 1, however, the result should be 2.
```

Rewrite the multi-insert statement by merging the branches inserting into the same partition to one branch.

#### **CDPD-57574: Query execution fails due to NullPointerException in DagUtils.setupQuickStart**

Hive queries may fail at runtime with a NullPointerException while executing the Tez graph at DagUtils.setupQuickStart method. A part of the stack trace from the error is shown below:

```
ERROR : Failed to execute tez graph.
java.lang.NullPointerException: null at org.apache.hadoop.hive.
ql.exec.tez.DagUtils.setupQuickStart(DagUtils.java:1724)
```

The error usually occurs from an invalid plan that has cycles created by semi-join or map-join edges.

Disable some of the optimizations in order to remove the cycles from the plan. It is sufficient to set one of the following properties:

- set hive.tez.dynamic.partition.pruning=false;
- set hive.tez.dynamic.semijoin.reduction=false;
- set hive.auto.convert.join=false;

To avoid any undesired performance regressions, disable the optimization(s) only for specific queries and not globally (through Cloudera Manager) for the whole workload or cluster.

#### **CDPD-43107: SemanticException for INSERT INTO statement when Hive Cost-based Optimizer (CBO) is disabled**

If you are running the INSERT INTO query and hive.cbo.enable is set to "false", the query fails with a SemanticException. For example,

```
set hive.cbo.enable=false;

CREATE TABLE mytable (
```

```
id INT,
str STRING
);

INSERT INTO mytable (id, str) VALUES (1, 'a');
```

Output:  
 org.apache.hadoop.hive.ql.parse.SemanticException: 0:0 Expected 2 columns for insclause-0/default@mytable; select produces 1 columns. Error encountered near token 'a'

This issue occurs only when columns are specified for the target table in the query. INSERT INTO mytable values (1, 'a'); does not result in an exception.

Enable CBO and run the query.

#### CDPD-43957: HiveServer shuts down during replication due to high resource usage

During Hive replication (both bootstrap and incremental), you may notice that the HiveServer (HS2) shuts down periodically with the following error:

```
java.sql.SQLException: org.apache.hive.jdbc.ZooKeeperHiveClientException: Unable to read HiveServer2 configs from ZooKeeper
 at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.java:265)
 at org.apache.hive.jdbc.HiveDriver.connect(HiveDriver.java:107)
 at java.sql.DriverManager.getConnection(DriverManager.java:664)
 at java.sql.DriverManager.getConnection(DriverManager.java:247)
 at com.cloudera.enterprise.hive3qt.Hive3QueryTool$HiveOperation.execute(Hive3QueryTool.java:682)
 at com.cloudera.enterprise.hive3qt.Hive3QueryTool.main(Hive3QueryTool.java:935)
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Unable to read HiveServer2 configs from ZooKeeper
 at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:177)
 at org.apache.hive.jdbc.Utills.configureConnParamsFromZooKeeper(Utills.java:580)
 at org.apache.hive.jdbc.Utills.parseURL(Utills.java:391)
 at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.java:263)
 ... 5 more
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Tried all existing HiveServer2 uris from ZooKeeper.
 at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.getServerHosts(ZooKeeperHiveClientHelper.java:132)
 at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:172)
 ... 8 more
```

This issue occurs when you replicate at scale with an unbalanced cluster setup that has all the roles running on the Cloudera Manager host. As a result, Cloudera Manager ends up in a bottleneck situation because HS2 crashes, preventing further replication. The logs (/var/log/messages) indicate that there were a large number of Java processes running on the host, which exhausted the host's memory and triggered an Out Of Memory Killer process to stop HS2.

You can restart HS2 and reinitiate replication. However, the replication process may take a longer time to complete.



- Follow these recommendations when setting up your source and target clusters:
  - Ensure that no HS2, Hive metastore (HMS), DataNode, or NameNode roles are running on the Cloudera Manager host.
  - Ensure that you have multiple instances of HS2 and HMS roles on different nodes.
  - It is recommended that you have the HS2 and HMS roles on different nodes than the DataNodes or NameNodes.

These practices increase the possibility of Hive replication at scale completing successfully.

- Limit the scale of Hive replication by temporarily disabling replication policies.

#### **CDPD-40730: Parquet change can cause incompatibility**

Parquet files written by the parquet-mr library in CDP 7.1.9, where the schema contains a timestamp with no UTC conversion will not be compatible with older versions of Parquet readers. The effect is that the older versions will still consider these timestamps as they would require UTC conversions and will thus end up with a wrong result. You can encounter this problem only when you write Parquet-based tables using Hive, and tables have the non-default configuration `hive.parquet.write.int64.timestamp=true`.

None.

#### **CDPD-41274: HWC + Oozie issue: Could not open client transport with JDBC Uri**

Currently only Spark cluster mode is supported in the Oozie Spark Action with Hive Warehouse Connector (HWC).

Use Spark action in cluster mode.

```
<spark xmlns="uri:oozie:spark-action:1.0">
 ...
 <mode>cluster</mode>
 ...
</spark>
```

#### **CDPD-26556 After an upgrade, querying a CTAS table under certain conditions might throw an exception**

If you upgrade your Hive cluster from CDH 6 to CDP 7, create a CTAS table in the CDP cluster from a table you upgraded from CDH, you might see the following exception when you query the new table:

```
class org.apache.hadoop.io.IntWritable cannot be cast to class o
rg.apache.hadoop.hive.serde2.objectinspector.StandardUnionObject
Inspector$StandardUnion
```

This issue involves CDH-based tables having columns of complex types ARRAY, MAP, and STRUCT.

#### **CDPD-45370: Cartesian join fails with an empty table**

When the configuration `hive.tez.cartesian-product.enabled` is set to true, query fails if one of the tables involved in a Cartesian join has zero rows. This results in an `IllegalArgumentException` with the message `Number of items is 0. Should be positive`.

To address this issue, set `hive.tez.cartesian-product.enabled` to false before running the query.

Apache Jira: [\[HIVE-26582\]](#)

#### **CDPD-23506: OutOfMemoryError in LLAP**

Long running spark-shell applications can leave sessions in interactive Hiveserver2 until the Spark application finishes (user exists from spark-shell), causing memory pressure in case of a high number of queries in the same shell (1000+).

You must close spark-shell so that sessions are closed. Add the owner of the database or the tables as a user with read or read/write access to the tables directly.

#### CDPD-23041: DROP TABLE on a table having an index does not work

If you migrate a Hive table to CDP having an index, DROP TABLE does not drop the table. Hive no longer supports indexes ([HIVE-18448](#)). A foreign key constraint on the indexed table prevents dropping the table. Attempting to drop such a table results in the following error:

```
java.sql.BatchUpdateException: Cannot delete or update a parent
row: a foreign key constraint fails ("hive"."IDXS", CONSTRAINT "
IDXS_FK1" FOREIGN KEY ("ORIG_TBL_ID") REFERENCES "TBLS" ("TBL_ID"
))
```

There are two workarounds:

- Drop the foreign key "IDXS\_FK1" on the "IDXS" table within the metastore. You can also manually drop indexes, but do not cascade any drops because the IDXS table includes references to "TBLS".
- Launch an older version of Hive, such as Hive 2.3 that includes IDXS in the DDL, and then drop the indexes as described in [Language Manual Indexing](#).

Apache Issue: [Hive-24815](#)

#### CDPD-17766: Queries fail when using spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive.

Do not specify spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive remotely.

Workaround: specify principal=hive.server2.authentication.kerberos.principal as shown in the following syntax:

```
jdbc:hive://<host>:<port>/<dbName>;principal=hive.server2.authen
tication.kerberos.principal;<otherSessionConfs>?<hiveConfs>#<hiv
eVars>
```

#### CDPD-13636: Hive job fails with OutOfMemory exception in the Azure DE cluster

Set the parameter hive.optimize.sort.dynamic.partition.threshold=0. Add this parameter in Cloudera Manager (Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml)

#### CDPD-10848: HiveServer Web UI displays incorrect data

If you enabled auto-TLS for TLS encryption, the HiveServer2 Web UI does not display the correct data in the following tables: Active Sessions, Open Queries, Last Max n Closed Queries

#### Known Issues in Hue

Learn about the known issues in Hue, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

#### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7

There are no new known issues identified in this release.

#### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7

##### CDPD-58978: Batch query execution using Hue fails with Kerberos error

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark queries in a batch mode launches a shell script Oozie job from Hue and this is not supported on a secure cluster.

There is no workaround. You can submit the queries individually.

##### OPSAPS-69659: Hue service fails on restart with "Unable to find pycpg2 2.5.4" error

Hue service fails to restart and you see the following error: Unable to find psycopg2 2.5.4. This could be because you have installed Python in a non-default location and Hue is unable to locate the psycopg2 PostgreSQL database adapter.

You must specify the path where you have installed Python in the PYTHONPATH property in the Hue Advanced Configuration Snippet using Cloudera Manager.

1. Log in to Cloudera Manager as an Administrator.
2. Go to Clusters Hue Configurations and add the following key and value in the Hue Service Environment Advanced Configuration Snippet (Safety Valve) field:

Key: PYTHONPATH

Value: `[***PYTHON-PATH***]`

Replace `[***PYTHON-PATH***]` with the actual location where you have installed Python. For example, `/opt/cloudera/parcels/CDH/lib/hue/build/env/lib/python3.8/site-packages`

3. Click Save Changes.
4. Restart the Hue service.

#### **CDPD-59677: Unable to view Phoenix tables on the left assist in Hue**

On clusters secured with Knox, you may not be able to see Phoenix tables on the left assist that are present under the default database (that is, an empty("") database).

None.

#### **CDPD-58142: A query is not pre-populated in the Hue editor after clicking on the “Re Execute” button**

When you click Re Execute to rerun a query from the Job Browser Queries Query Details page, the query does not get populated on the Hue editor, as expected.

None.

#### **CDPD-54376: Clicking the home button on the File Browser page redirects to HDFS user directory**

When you are previewing a file on any supported filesystem, such as S3 or ABFS, and you click on the Home button, you are redirected to the HDFS user home directory instead of the user home directory on the said filesystem.

None.

#### **CDPD-41306: pip3.8 freeze command does not work and results into an error**

You see the `/usr/lib/hue/build/env/bin/python: No such file or directory` error when you run the following command:

```
build/env/bin/pip3.8 freeze
```

Run the freeze command as follows by specifying the paths of python3.8 and pip3.8:

```
/opt/cloudera/parcels/CDH/lib/hue/build/env/bin/python3.8 /opt/cloudera/parcels/CDH/lib/hue/build/env/bin/pip3.8 freeze
```

#### **CDPD-43293: Unable to import Impala table using Importer**

Creating Impala tables using the Hue Importer may fail.

If you have both Hive and Impala services installed on your cluster, then you can import the table using by selecting the Hive dialect from Tables Sources .

If only Impala service is installed on your cluster, then go to Cloudera Manager Clusters Hue Configurations and add the following line in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue\_safety\_valve.ini field:

```
[beeswax]
max_number_of_sessions=1
```

**DWX-8602: Unable to import a large CSV file from the local workstation**

You may see an error message while importing a CSV file into Hue from your workstation, stating that you cannot import files of size more than 200 KB.

Upload the file to S3 or ABFS and then import them into Hue using the Importer.

**CDPD-39330: Unable to use the pip command in CDP**

You may not be able to use the pip command in CDP 7.1.7 or higher and may see the following error when using pip in a command: “ImportError: cannot import name chardet”.

Follow the steps listed on [Unable to use pip command in CDP](#).

**CDPD-24294: Hue uses the unsafe-inline directive in its Content Security Policy (CSP) header**

Hue 4 web interface uses the unsafe-inline directive in its CSP header. As a result, the application server does not set the CSP header in its HTTP responses, and therefore does not benefit from the additional protection against potential cross-site scripting issues and other modern application vulnerabilities which a properly configured CSP may provide. This could lead to application vulnerability.

Cloudera recommends deploying additional security measures such as a firewall within the Hue server to control allowed connections, and SSO-based authentications mechanisms such as LDAP or SAML.

**OPSAPS-58927: Connection failed error when accessing the Search app (Solr) from Hue**

If you are using Solr with Hue to generate interactive dashboards and for indexing data, and if you have deployed two Solr services on your cluster and selected the second one as a dependency for Hue, then Cloudera Manager assigns the hostname of the first Solr service and the port number of the second Solr service generating an incorrect Solr URL in the search section of the hue.ini file. As a result, you may see a “Connection failed” error when you try to access the Search app from the Hue web UI.

1. Log into Cloudera Manager as an Administrator.
2. Go to Clusters Hue service Configuration and add the following lines in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue\_safety\_valve.ini field:

```
[search]
URL of the Solr Server
solr_url=http://[***HOSTNAME***]:[***PORT***]/solr/
```

For example:

```
solr_url=http://solr2:4567/solr/
```

3. Click Save Changes.
4. Restart the Hue service.

**CDPD-41136: Importing files from the local workstation is disabled by default**

Cloudera has disabled the functionality to import files from your local workstation into Hue because it may cause errors. You may not see the Local File option in the Type drop-down menu on the **Importer** page by default.

You can enable the functionality to import files from your local workstation by specifying the following parameter in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue\_safety\_valve.ini field in Hue configurations in Cloudera Manager:

```
[indexer]
enable_direct_upload=true
```

**Invalid S3 URI error while accessing S3 bucket**

The Hue Load Balancer merges the double slashes (//) in the S3 URI into a single slash (/) so that the URI prefix `"/filebrowser/view=S3A:/"` is changed to `"/filebrowser/view=S3A:/"`. This results in an error when you try to access the S3 buckets from the Hue File Browser through the port 8889.

The Hue web UI displays the following error: “Unknown error occurred”.

The Hue server logs record the “ValueError: Invalid S3 URI: S3A” error.

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.

You should be able to load the S3 browser from both 8888 and 8889 ports.

Alternatively, you can use the Hue server port 8888 instead of the load balancer port 8889 to resolve this issue.

### Error while rerunning Oozie workflow

You may see an error such as the following while rerunning an already executed and finished Oozie workflow through the Hue web interface: `E0504: App directory [hdfs://cdh/user/hue/oozie/workspaces/hue-oozie-1571929263.84] does not exist`.

To resolve this issue, add the following property in the Hue Load Balancer Advanced Configuration Snippet:

1. Sign in to Cloudera Manager as an administrator.
2. Go to `Clusters Hue service Configurations Load Balancer` and search for the Load Balancer Advanced Configuration Snippet (Safety Valve) for `httpd.conf` field.
3. Specify `MergeSlashes OFF` in the Load Balancer Advanced Configuration Snippet (Safety Valve) for `httpd.conf` field.
4. Click `Save Changes`.
5. Restart the Hue Load Balancer.

### CDPD-17956, CDPD-14167: Setting idle session timeout for Hue does not work when the cluster is secured using Knox SSO

If Hue is configured with `desktop.auth.backend.KnoxSpnegoDjangoBackend` as the Authentication Backend, then the automatic idle session logout that is set by configuring the `idle_session_timeout` property does not take effect. You may also see 404 error while accessing Hue from the Knox UI when the `idle_session_timeout` property is not set to -1.

None

### Query history displays "Result Expired" message

You see the "Result Expired" message under the Query History column on the **Queries** tab for queries which were run back to back. This is a known behaviour.

None.

### CDPD-43956: Manually replace UUID when importing Oozie workflows containing sub-workflows

When importing Oozie workflows that contain sub-workflows, you must replace all the UUID entries with unique new entries. If you change the UUID of a sub-workflow, you must update that reference in the parent workflow to avoid circular dependencies.

### Hue limitation after upgrading from CDH to CDP Private Cloud Base

The `hive.server2.parallel.ops.in.session` configuration property changes from `TRUE` to `FALSE` after upgrading from CDH to CDP Private Cloud Base. Current versions of Hue are compatible with this property change; however, if you still would like to use an earlier version of Hue that was not compatible with this property being `FALSE` and shared a single JDBC connection to issue queries concurrently, the connection will no longer work after upgrading.

### Hue UI is blank upon login after upgrading to CDP 7.1.7 from CDH 6

If your cluster was secured using Knox, and if you have upgraded from CDH 6 to CDP 7.1.7, then you may see a blank Hue screen. This could happen because the `knox_proxyhosts` parameter is newly introduced in CDP, and it is possible that this parameter is not configured in Cloudera Manager under Hue configuration.

Specify the host on which you have installed Knox in the Hue Knox Proxy Hosts configuration as follows:

1. Log in to Cloudera Manager as an Administrator.
2. Obtain the host name of the Knox Gateway by going to Clusters Knox service Instances .
3. Go to Clusters Hue service Configuration and search for the Knox Proxy Hosts field.



**Note:** Cloudera Manager displays the following warning if the Knox Proxy Hosts field is empty when Knox Gateway is enabled on the CDP cluster: The parameter `knox_proxyhosts` cannot be empty. This can happen if there are no Knox Gateways. Please set the `knox_proxyhosts` to the list of hosts that have Knox Gateways.

4. Specify the Knox Gateway hostname in the Knox Proxy Hosts field.
5. Click Save Changes and restart the Hue service.

### Unsupported features

#### CDPD-59595: Spark SQL does not work with all Livy servers that are configured for High Availability

SparkSQL support in Hue with Livy servers in HA mode is not supported. Hue does not automatically connect to one of the Livy servers. You must specify the Livy server in the Hue Advanced Configuration Snippet as follows:

```
[desktop]
[spark]
livy_server_url=http(s)://[***LIVY-FOR-SPARK3-SERVER-HOST***]:
[***LIVY-FOR-SPARK3-SERVER-PORT***]
```

Moreover, you may see the following error in Hue when you submit a SparkSQL query: Expecting value: line 2 column 1 (char 1). This happens when the Livy server does not respond to the request from Hue.

Specify all different Livy servers in the `livy_server_url` property one at a time and use the one which does not cause the issue.

#### CDPD-18491: PySpark and SparkSQL are not supported with Livy in Hue

Hue does not support configuring and using PySpark and SparkSQL with Livy in CDP Private Cloud Base.

### Importing and exporting Oozie workflows across clusters and between different CDH versions is not supported

You can export Oozie workflows, schedules, and bundles from Hue and import them only within the same cluster if the cluster is unchanged. You can migrate bundle and coordinator jobs with their workflows only if their arguments have not changed between the old and the new cluster. For example, hostnames, NameNode, Resource Manager names, YARN queue names, and all the other parameters defined in the `workflow.xml` and `job.properties` files.

Using the import-export feature to migrate data between clusters is not recommended. To migrate data between different versions of CDH, for example, from CDH 5 to CDP 7, you must take the

dump of the Hue database on the old cluster, restore it on the new cluster, and set up the database in the new environment. Also, the authentication method on the old and the new cluster should be the same because the Oozie workflows are tied to a user ID, and the exact user ID needs to be present in the new environment so that when a user logs into Hue, they can access their respective workflows.



**Note:** Migrating Oozie workflows from HDP clusters is not supported.

### Known Issues in Iceberg

Learn about the known issues in Iceberg, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7

There are no known issues identified in earlier releases.

### Known Issues in Impala

Learn about the known issues in Impala, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7

#### CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

#### CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the `--tls_ciphersuites` startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.

Ensure that the list in the `--tls_ciphersuites` startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

**Apache Jira:** [IMPALA-14625](#)

**CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.**

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

**CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work**

If you try to run impala-shell on RHEL 9 by setting the default python executable available in PATH to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, IMPALA\_PYTHON\_EXECUTABLE=python3.

**CDPD-42958: After upgrading the CDH 7.1.9 from CDH 6.x, under certain conditions you cannot insert data into a table**

Under the following conditions, after upgrading from CDH 6.x to CDH 7.1.9 you cannot insert data into a table from Impala:

- On CDH 6.x, you created a database with Impala in a user specified HDFS location.
- Using Hive, you then created a table in the database.

Under these conditions, the database and table are stored in the user-specified HDFS directory. After upgrading, the HDFS directory of the table is read-only for Impala. Consequently, from Impala you cannot insert new data into the table because Impala does not have write permission on the HDFS directory.

Workaround: To resolve this issue, use either one of the following workarounds:

- Using the Ranger Web UI, in the policy repository cm\_hdfs, grant the user 'impala' write permission on the directory where the table resides.
- Enter the following command to grant write permission to user 'impala' on the HDFS directory where the table resides.

```
hdfs dfs -setfacl -m default:user:impala:rwX <HDFS directory>
```

**CDPD-80019: Impala fails to start due to large catalog-update topic**

Impala coordinators fail to start if the catalog-update statestore topic exceeds 2GB, hitting the Thrift 0.16.0 message size limit (thrift\_rpc\_max\_message\_size). This prevents the coordinator from processing the update.

- Enable local catalog mode: --use\_local\_catalog=true
- In Cloudera Manager, go to Impala Configuration search for "Catalog Server Command Line Argument Advanced Configuration Snippet (Safety Valve)" and click 
- Add property name and value:

```
Name : --catalog_topic_mode
Value: minimal
```

- Save changes.

Apache Jira: [IMPALA-12350](#)

**Impala cannot update table if the 'external.table.purge' property is not set to true**

Impala cannot update a table using DDL statements if the 'external.table.purge' property is FALSE. ALTER TABLE statements return success with no changes to the table.

ALTER TABLE statements should be issued twice if "external.table.purge" was FALSE initially.



**Note:** For Iceberg tables, Impala users should always use CREATE TABLE since the Impala CREATE TABLE statement sets the flag to TRUE. However, Impala CREATE EXTERNAL TABLE sets the flag to FALSE.

**Impala's known limitation when querying compacted tables**



When the compaction process deletes the files for a table from the underlying HDFS location, the Impala service does not detect the changes as the compactions does not allocate new write ids. When the same table is queried from Impala it throws a 'File does not exist' exception that looks something like this:

```
Query Status: Disk I/O error on <node>:22000: Failed to open HDFS file hdfs://nameservice1/warehouse/tablespace/managed/hive/<database>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException: File does not exist: /warehouse/tablespace/managed/hive/<database>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

**CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.**

You must use a single Knox node to access Impala UI.

**Impala api calls via 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-21828: Multiple permission assignment through grant is not working**

None

**Problem configuring masking on tables using Ranger**

The following Knowledge Base article describes the behavior when we configure masking on tables using Ranger. This configuration works for Hive, but breaks queries in some scenarios for Impala.

For a workaround, see the following Knowledge Base article: [ERROR: "AnalysisException: No matching function with signature: mask\(FLOAT\)" when Impala jobs fail with the following error with signature: mask\(FLOAT\)](#)

**IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS**

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm\_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXXXX"

**IMPALA-532: Impala should tolerate bad locale settings**

If the LC\_\* environment variables specify an unsupported locale, Impala does not start.

Add LC\_ALL="C" to the environment settings for both the Impala daemon and the Statestore daemon.

**Avro Scanner fails to parse some schemas**

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

## IMPALA-691: Process mem limit does not account for the JVM's memory usage

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the impalad daemon.

To monitor overall memory usage, use the `top` command, or add the memory figures in the Impala web UI `/memz` tab to JVM memory usage shown on the `/metrics` tab.

## IMPALA-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Remove trailing semicolon from the Avro schema.

### IMPALA-1652: Incorrect results with basic predicate on CHAR typed column

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the `RPAD()` function to blank-pad literals compared with `CHAR` columns to the expected length.

### IMPALA-1821: Casting scenarios with invalid/inconsistent results

Using a `CAST()` function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

### IMPALA-2005: A failed CTAS does not drop the table if the insert fails

If a `CREATE TABLE AS SELECT` operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

## Drop the new table manually after a failed CREATE TABLE AS SELECT

## IMPALA-2422: % escaping does not work correctly when occurs at the end in a LIKE clause

If the final character in the RHS argument of a LIKE operator is an escaped \% character, it does not match a % final character of the LHS argument.

None

## IMPALA-2603: Crash: impala::Coordinator::ValidateCollectionSlots

A query could encounter a serious error if includes multiple nested levels of INNER JOIN clauses involving subqueries.

None

### IMPALA-3094: Incorrect result due to constant evaluation in query with outer join

An OUTER JOIN query could omit some expected result rows due to a constant such as FALSE in another join clause. For example:

```
explain SELECT 1 FROM alltypeptyny 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 |
```

[illegible]

### IMPALA-3509: Breakpad minidumps can be very large when the thread count is high

The size of the breakpoint minidump files grows linearly with the number of threads. By default, each thread adds 8 KB to the minidump size. Minidump files could consume significant disk space when the daemons have a high number of threads.

Add `-minidump_size_limit_hint_kb=size` to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then 2 KB per thread after that.) The minidump file can still grow larger than the "hinted" size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

## IMPALA-4978: Impala requires FQDN from hostname command on Kerberized clusters

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the `gethostname()` system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a Kerberized cluster.

Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname` only returns the short name, pass the command-line flag `##hostname=FULLY_QUALIFIED_DOMAIN_NAME` in the startup options of all Impala-related daemons.

## IMPALA-6671: Metadata operations block read-only operations on unrelated tables

Metadata operations that change the state of a table, like `COMPUTE STATS` or `ALTER RECOVER PARTITIONS`, may delay metadata propagation of unrelated unloaded tables triggered by statements like `DESCRIBE` or `SELECT` queries.

None

## IMPALA-7072: Impala does not support Heimdal Kerberos

None

**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.

## 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.

### Known Issues in Kafka

Learn about the known issues in Kafka, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7

#### **CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress**

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

#### **OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners**

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL\_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker) in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Save Changes > Restart SMM.

#### **The `offsets.topic.replication.factor` property must be less than or equal to the number of live brokers**

The `offsets.topic.replication.factor` broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a `GROUP_COORDINATOR_NOT_AVAILABLE` error until the cluster size meets this replication factor requirement.

None

#### **Requests fail when sending to a nonexistent topic with `auto.create.topics.enable` set to true**

The first few produce requests fail when sending to a nonexistent topic with `auto.create.topics.enable` set to true.

Increase the number of retries in the producer configuration setting `retries`.

#### **Performance degradation when SSL is enabled**

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with `ssl.secure.random.implementation = SHA1PRNG`. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

#### **OPSAPS-43236: Kafka garbage collection logs are written to the process directory**

By default Kafka garbage collection logs are written to the agent process directory. Changing the default path for these log files is currently unsupported.

None

#### **RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure**

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transactional method because we are in an error state
 at org.apache.kafka.clients.producer.internals.TransactionManager.maybeFailWithError(TransactionManager.java:1125)
 at org.apache.kafka.clients.producer.internals.TransactionManager.maybeAddPartition(TransactionManager.java:442)
 at org.apache.kafka.clients.producer.KafkaProducer.doSend(KafkaProducer.java:1000)
 at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:914)
 at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:800)
 .
 .
 .
 Caused by: org.apache.kafka.common.errors.ClusterAuthorizationException: Cluster authorization failed.
```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting `enable.idempotence` to `false`.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

#### **CDPD-49304: AvroConverter does not support composite default values**

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

#### **DBZ-4990: The Debezium Db2 Source connector does not support schema evolution**

The Debezium Db2 Source connector does not support the evolution (updates) of schemas. In addition, schema change events are not emitted to the schema change topic if there is a change in the schema of a table that is in capture mode. For more information, see [DBZ-4990](#).

None.

#### **CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression**

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy compression. The HDFS Stateless Sink connector is only affected if the `Compression Codec` or `Compression Codec` for Parquet properties are set to `SNAPPY`.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: org.apache.hadoop.util.NativeCodeLoader.buildSupportsSnappy()
```

```
Failed to write to HDFS due to java.lang.RuntimeException: native snappy library not available: this version of libhadoop was built without snappy support.
```

Download and deploy missing libraries.



**Important:** Ensure that you complete steps 1-11 on all Kafka Connect hosts. Additionally, ensure that the advanced configuration snippet in step 12 is configured for all Kafka Connect role instances.

1. Create the /opt/nativelibs directory.

```
mkdir /opt/nativelibs
```

2. Change the owner to kafka.

```
chown kafka:kafka /opt/nativelibs
```

3. Locate the directory containing the Hadoop native libraries and copy its contents to the directory you created.

```
cp /opt/cloudera/parcels/CDH/lib/hadoop/lib/native/* /opt/nativelibs
```

4. Verify that libsnpappy.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**

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 Key Trustee Server**

Learn about the known issues in Key Trustee Server (KTS), the impact or changes to the functionality, and the workaround.

##### **KTS setup on FIPS cluster gets stuck if KTS is in HA**

Key Trustee Server setup gets stuck if KTS is in HA for Centos 7.9 in FIPS mode.

##### **Create Date is not fetched from Key Trustee Server**

During key migration from KTS to Ranger KMS, create date is not fetched from KTS.

#### **Known Issues in Knox**

Learn about the known issues in Knox, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7**

##### **CDPD-27801: Knox is missing HSTS header for HTTP 404 responses**

Knox does not support the HTTP Strict-Transport-Security response header (HSTS) in some HTTP 404 responses.

None.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7**

##### **CDPD-61527: On FIPS cluster, Knox-Impala connection is failing with SSL error code 5**

The Knox-Impala connection is failing with SSL error code 5 on FIPS clusters. This prevents accessing Impala through Knox.

#### **Known Issues in Kudu**

Learn about the known issues in Kudu, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7**

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7

**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.

Workaround: Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

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.

Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

### Unsupported features

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 in Cloudera Runtime 7.1.9 SP1 CHF 7.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7

The following known issues were identified in this release:

No new known issues in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7

No known issues identified in Livy.

### Known issues in Navigator Encrypt

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround.

#### Failed to start navencrypt-mount.service

After installing NavEncrypt on SLES, if the command "systemctl status navencrypt-mount" fails with the error: "Failed to start navencrypt-mount.service: Unit navencrypt-mount.service failed to load: No such file or directory", the systemd files need to be reloaded.

Run the command:

```
(sudo) systemctl daemon-reload
```

#### Crash in Navigator Encrypt kernel module if --profile-file= option is used with ACLs

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 Ozone

Learn about the known issues in Ozone the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7

The following known issues were identified in this release:



**CDPD-81569: Recon is not reflecting correct metrics for Volumes, Buckets, Keys, and Pending Key Deletions**

Recon is displaying data for container and space utilisation only. The metrics for Volumes, Buckets, Keys, and Pending Key Deletions are not being updated and displayed correctly.

Apache Jira: [HDDS-12607](#)

**CDPD-84518: Expired secret key can abort leader Ozone Manager startup**

Ozone Manager fails to startup due to NPE exception because of expired SCM managed secret key.

Apache Jira: [HDDS-13234](#)

**Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7****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: 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.

**The Prometheus binary is not available in CDP Private Cloud Base 7.1.9 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.

**After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and if Ozone is in the Non-HA environment, an exception message is observed during the finalization of the Ozone upgrade.**

During the finalization of the upgrade, a `ClassNotFoundException` message for `org.cloudera.log4j.redactor.RedactorAppender` class may be displayed. The error message can be ignored because the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

**CDPD-68951: In CDP Private Cloud Base 7.1.9 CHF2 version and lower, the command `ozone sh key list <bucket_path>` displays the `isFile` flag in a key's metadata as `false` even when the key is a file. This issue is rectified in CDP Private Cloud Base 7.1.9 CHF3. However, the pre-existing (pre-upgrade) key's metadata cannot be changed.**

None

**CDPD-67771: When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.**

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(OzoneManager
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.lambda
$1(OzoneManagerStateMachine.java:363) at
java.base/java.util.concurrent.CompletableFuture
$AsyncSupply.run(CompletableFuture.java:1700) at java.base/
java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1128)
at java.base/java.util.concurrent.ThreadPoolExecutor
$Worker.run(ThreadPoolExecutor.java:628) at java.base/
java.lang.Thread.run(Thread.java:834)
```

This occurs because S3A uses multipart upload to commit job results in a batch. The staging committer's replace mode deletes the target directory before completing Memory Protection Unit (MPU). The problem is that File System Optimization (FSO) does not create intermediate directories during MPU, it does only for regular file/dir/key requests.

Use `fs.s3a.committer.name=magic` for affected versions.

**CDPD-64398: In SCM with non-HA configuration, Secret key manager does not initialise. Hence, the startup of OM and DataNode fails because it cannot get a secret key. This key is used when security is enabled and used in Block token and container token verification while communicating between the Ozone client, OM, and Datanode.**

This issue does not occur in SCM with HA configuration.

You must force exit from the safe mode for SCM. This triggers the initialization of the Secret key manager. Below are the options:

- exit safe mode manually whenever SCM is started
- disable safe mode by setting `hdds.scm.safemode.enabled=false` in safety valve for SCM configuration

Impact of disabling safe mode and the purpose of safe mode:

- Write should not fail once the cluster is out of safe mode
- Read of existing data should not fail after the cluster is out of safe mode
- Unnecessary re-replication should be avoided during cluster restart

So disabling safe mode does not have any major impact on Ozone.

**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.

**CDPD-52412: keyManagerImpl#listStatus exceeds the maximum RPC length.**

To reduce the payload, you must increase the part size.

**OPSAPS-68159: If you did not deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Error when distributing parcel to host, Ignoring non-compliant parcel manifest error is displayed after Cloudera Manager is upgraded to 7.11.3.**

If you encounter the error, perform the following steps:

1. Deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.11.3.
2. Restart the cluster with a delay of 10 minutes.
3. Continue to upgrade CDH 7.1.8 to CDH 7.1.9.

**OPSAPS-68159: If you did not deactivate the Ozone parcel 718.2.x on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Ozone roles during the CDH 7.1.8 upgrade to CDH 7.1.9.**

If you encounter the error, perform the following steps:

1. Deactivate the Ozone parcel 718.2.x.
2. Restart the Ozone service.
3. Perform Finalize Upgrade for the Ozone service.

Step result: The Ozone roles will be displayed in green.

**CDPD-60989: The packaging version for Apache Ozone points to the 1.2.0 older version. This is a version string problem and not a packaging issue. The version of the Apache Ozone binary is closest to 1.4.0.**

None. This only affects the JAR versioning.

**CDPD-60366: Native library loader fails when system property native.lib.tmp.dir is not set. It fails because the library is copied to / instead of the cwd.**

Setting native.lib.tmp.dir to a certain path like /tmp should solve the issue. To set this system property, you must add -Dnative.lib.tmp.dir=/tmp to ozone\_java\_opts in Cloudera Manager configuration of the Ozone cluster.

**CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.**

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

**CDPD-60598: DataNode can consume up to the JVM heap size worth of direct memory buffers under specific failure scenarios. This can lead to a DataNode crash. This is due to Netty allocating and not freeing the direct memory buffers.**

Restarting DataNode will resolve this issue.

**CDPD-60466: The Ozone client is missing in the Cloudera Manager's Service Monitor node. This causes Ozone Canary's health check to fail.**

Ensure you install Ozone CLI on the same node as Cloudera Manager's Service Monitor.

**CDPD-60012: The Ozone SCM may be stuck in safe mode after upgrade to 7.1.9. This is due to a bug in how the SCM accounts for open storage containers.**

The workaround is to temporarily set the Ozone configuration hdds.scm.safemode.threshold.pct to a lower value like 0.90 and restart the SCM.

**CDPD-59679: In some scenarios, excess UNHEALTHY replicas of EC containers may not be removed from the cluster.**

None

**CDPD-50116: Topology-aware reads can provide better performance by routing applications to the closest replica of a file, if available, on the same physical node or same rack. However this feature is disabled by default in CDP 7.1.9.**

This feature can be enabled by setting ozone.network.topology.aware.read to true using Cloudera Manager safety valve.

**CDPD-57165: The DataNode disk usage thread may abort due to Ratis log tailing.**

Use `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory` to calculate disk usage. Configure in the `ozone-site.xml` property: `hdds.datanode.du.factory` value: `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory`

**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.

**CDPD-55513: Hive external table replication policy does not migrate Hive tables in HDFS storage to Ozone.**

To replicate data from HDFS to Ozone, see [Migrating your data from HDFS to Ozone](#). You can use HMS Mirror to replicate the metadata.

**OPSAPS-67757: Hive external tables in Ozone storage cannot be replicated using Hive external table replication policies.**

To replicate the Hive external tables' data, consider using DistCp. To replicate the metadata of Hive external tables, consider using HMS Mirror.

**Remove bucket recursively using `rb --force` command from AWS S3 cannot work for FSO buckets.**

Use the Ozone shell command `ozone sh bucket delete -r [***BUCKET ADDRESS***]`

**CDPD-59126: Info log displays "noexec permission on /tmp/liborg\_apache\_ratis\_thirdparty\_netty\_transport\_native\_epoll\_x86" on client while executing command with noexec on /tmp.**

To suppress Info log related to `liborg_apache_ratis_thirdparty_netty_transport_native_epoll_x86` library: Export `OZONE_OPTS` environment variable on the client terminal by running the command `export OZONE_OPTS="-Dorg.apache.ratis.thirdparty.io.netty.native.workdir=/var/tmp $OZONE_OPTS"`

**OPSAPS-67650: Ozone uses RocksDB as a library to persist metadata locally.**

By default, RocksDB places certain executables in `/tmp`, and thus encounters errors when `/tmp` is mounted with `noexec`.

The workaround is to configure RocksDB to put executables at another location. On a PhatCat node, the steps are:

1. Go to Cloudera Manager UI OZONE Configuration.
2. Find Ozone Service Environment Advanced Configuration Snippet (Safety Valve) and set the following environment variable: `ROCKSDB_SHAREDLIB_DIR=/VAR/TMP`
3. Restart Ozone.

**CDPD-49137: Ozone Manager Kerberos token expires for SCM communication and OM does not log in again.**

Sometimes, OM's Kerberos token is not updated and it stops to communicate with SCM. When this occurs, writes start failing.

Restart OM or set the safety valve `hadoop.kerberos.keytab.login.autorenewal.enabled = true`

#### **CDPD-56684: Keys get deleted when you do not have permission on volume**

When a volume is deleted, it recursively deletes the buckets and keys inside it and only then deletes the volume. The volume delete ACL check is done only in the end, due to which you may end up deleting all the data inside the volume without having delete permission on the volume.



**Note:** There was no bucket/key permission available which allowed the user to delete them recursively.

#### **CDPD-50610: Large file uploads are slow with OPEN and stream data approach**

Hue file browser uses the append operation for large files. This API is not supported by Ozone in 7.1.9 and therefore large file uploads can be slow or timeout from the browser.

Use native Ozone client to upload large files instead of the Hue file browser.

#### **OPSAPS-64097: Ozone service restart failed at SCM**

Stopping SCM service using Cloudera Manager can sometimes timeout and need a retry. The Cloudera Manager API waits for 90 seconds which is not sufficient under certain circumstances.

Retry the shutdown using Cloudera Manager if the SCM still shows as running after a refresh of the service status.

#### **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-filesystem-hadoop3.jar`, it is likely due to the presence of the user `hdfs` on the node prior to CDP parcel activation. As a result, the Cloudera Manager agent skips the initialization related to HDFS, leading to the non-creation of the `/var/lib/hadoop-hdfs` directory.

Create the directory `“/var/lib/hadoop-hdfs”` on all nodes followed by the deactivation and activation of the CDP parcel (deactivate and activate the Ozone parcel instead, in case Ozone parcel is used).

#### **CDPD-50447: When SCM High Availability is enabled, each of the SCM web UIs report the host of the web UI as the leader of HA, and the other two as followers. This gives wrong information**

Correct output is available by running the `ozone admin scm roles --service-id=<ID>` command.

#### **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](#).

**OPSAPS-62327: In an Ozone cluster without any gateway roles, Ozone is unable to deploy client configurations and displays the ConfigGenException error.**

You must add the Ozone gateway roles to the cluster.

**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-36389: The configurations `datanodes.involved.max.percentage.per.iteration` and `size.moved.max.per.iteration` are meant to limit the max number of DataNodes that is involved and maximum size that can move in an iteration. This bug causes balancer to stop an iteration when it is two DataNodes or 1 Container size (5GB) away from hitting these limits. However, these DataNodes can again be considered for balancing in the next iteration. This means the cluster is balanced after enough iterations, albeit a bit slowly. This bug is apparent in small clusters of around four DataNodes where the**

**DataNodes can be either the source or target for a lot of moves but the iteration gets stopped when three DataNodes have been involved. It takes a higher number of iterations to eventually balance this cluster. While this is a performance issue, it does not prevent balancer from ultimately balancing the cluster.**

Increase the speed for balancing by decreasing the interval between each iteration using the configuration `balancing.iteration.interval`. The value of this configuration must be greater than `hdds.datanode.du.refresh.period`. `size.moved.max.per.iteration` can be increased to allow more data to move in one iteration.

**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-35141: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2> (state=08S01,code=40000) java.sql.SQLException: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2>. This issue does not occur if the source and target buckets are different in Hive queries. Currently, copying across the same bucket is only supported.**

Avoid different buckets in source and target path. Ensure that `hive.exec.stagingdir` is unset or set to the same bucket as the warehouse so that staging directories are created in the same bucket as the warehouse.

**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-41184: With Legacy buckets, FileSystem Optimised is not interoperating with the Ozone shell command. The cause is that the directory key entry in the DB KeyTable is stored as "dir1/" with trailing slash. But while performing the described operation, Ozone shell (o3://) is normalizing the given path and removing the trailing slash "/" from it. That results in KEY\_NOT\_FOUND exception.**

There are three workarounds:

- Use `FileSystem` API to Delete the Directories rather than `Shell-Command` API.
- Use `FSO` buckets instead of `Legacy Buckets`. As in `FSO`, you can create `Intermediate Directories` and `Delete Directories` using the `Ozone shell` commands.
- Disable and set the configuration `ozone.om.enable.filesystem.pathsflag` to `false` in order to delete the directories. This is generally not a preferred workaround because the cluster must be restarted to pick up the new changes.

**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-12542: Mount of Ozone filesystem with the help of FUSE fails.**

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-42897: EC writes are failing with "No enough datanodes to choose" after EC replication config is set globally.**

EC writes start failing when large number of pipelines are created as a result of multiple EC configs with different chunk sizes used to write keys.

If standard EC configs (that is, rs-3-2-1024k) are used to write keys, number of pipelines created per DataNode is limited to five and this issue is not seen with standard EC configs.

The recommendation is not to create too many random chunk sizes. It is configurable because, users can decide based on their workload. But do not to have separate chunk sizes for each file.

**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-40560: Filesystem Operations using Hadoop S3a connector on a `FILE_SYSTEM_OPTIMIZED` bucket fails. `org.apache.hadoop.ozone.om.exceptions.OMException: Unable to get file status: volume: s3v bucket: fso key: test/`**

Do not run Hadoop S3a commands on a `FILE_SYSTEM_OPTIMIZED` bucket. Use `OBJECT_STORE` bucket layouts.

**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.

**CDPD-43432: Ozone Service in fault state in DataNode - Long Running setup.**

Upgrade RocksDB to the latest version.

**OPSAPS-63999: In the newly installed cluster, the Finish upgrade option is clickable.**



None

**OPSAPS-64648: Failed to start Ozone node using Cloudera Manager if default log path /var/log/hadoop-ozone does not exist. If this path does not exist, any Ozone nodes(for example SCM or DataNode) restart will fail.**

Run the following command `sudo -u hdfs mkdir -p /var/log/hadoop-ozone` or replace hdfs with the user Ozone roles that are running.

**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.

### Technical Service Bulletins

**TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager**

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.

### Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

### Known issues in Ranger

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround.

### UnsupportedClassVersionError

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.

**CDPD-68806: The Revoke operation for users belonging to a group or role permission does not function as expected.**

List command is listing all the tables even when the user permission is revoked. And also the command does not add any deny policy to Ranger for that specific user.

For example:

user\_1 is the part of group\_1

user\_2 is the part of role\_1

Create a Ranger policy for group\_1 and role\_1 to access all the resources in HBase.

Try to revoke the permission for user\_1 and user\_2 using HBase shell global revoke command:

```
revoke 'user_1'
```

```
revoke 'user_2'
```

But there is no deny policy added for that user. And access is still granted to user\_1 and user\_2.

This behavior is currently not supported in HBase shell. Must be handled manually using the Ranger policy change.

#### **CDPD-68739: The revoke command does not work when using the HBase shell**

While using the HBase shell, running the revoke command does not cancel the user permission. Users are able to perform actions even after running the revoke command.

For example:

When a user is granted permission to a particular table using the HBase grant command.

```
hbase:003:0> grant 'user_1', 'RWCA', 'some_abc_table'
```

And then execute the global revoke command for that user:

```
hbase:005:0> revoke 'user_1'
```

Permission for that user 'user\_1' for the table 'some\_abc\_table' not got revoked.

None.

#### **CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.**

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

#### **CDPD-58860: After upgrading CDP-7.1.8 to CDP-7.1.9, cdp-proxy token missing from Knox-Ranger policy.**

As part of [OPSAPS-67480](#) in 719 default ranger policy is added from cdp-proxy-token topology, so that after a new installation of CDP-7.1.9, the Knox-Ranger policy includes cdp-proxy-token. However, upgrades do not add cdp-proxy-token to cm\_knox policies automatically.

Manually add cdp-proxy-token to the Knox policy, using Ranger Admin Web UI.

1. Log in to Cloudera Manager Ranger Admin Web UI, as a Ranger administrator.
2. On Ranger Admin Web UI Service Manager Resource Knox, click cm\_knox.
3. In Knox Policies, open the CDP Proxy UI, API and Token policy.
4. In Knox Topology\*, add cdp-proxy-token.
5. Click Save.
6. Restart Ranger.

#### **CDPD-58704: hadoop roll key / key delete command shows operation failed error when one KMS host is down, even when operation succeeds**

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered kms instances. If one KMS instance is stopped (not deleted), client gets 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.**

```
hbase:001:0> revoke 'hrt_11'
ERROR: org.apache.hadoop.hbase.coprocessor.CoprocessorException: HTTP 400 Error: processSecureRevokeRequest processing failed
 at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocessor.preRevoke(RangerAuthorizationCoprocessor.java:1309)
 at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocessor.preRevoke(RangerAuthorizationCoprocessor.java:1128)
 at org.apache.hadoop.hbase.master.MasterCoprocessorHost$162.call(MasterCoprocessorHost.java:1857)
 at org.apache.hadoop.hbase.master.MasterCoprocessorHost$162.call(MasterCoprocessorHost.java:1854)
 at org.apache.hadoop.hbase.coprocessor.CoprocessorHost$ObserverOperationWithoutResult.callObserver(CoprocessorHost.java:558)
 at org.apache.hadoop.hbase.coprocessor.CoprocessorHost.executeOperation(CoprocessorHost.java:631)
 at org.apache.hadoop.hbase.master.MasterCoprocessorHost.preRevoke(MasterCoprocessorHost.java:1854)
 at org.apache.hadoop.hbase.master.MasterRpcServices.revoke(MasterRpcServices.java:2740)
 at org.apache.hadoop.hbase.shaded.protobuf.generated.MasterProtos$MasterService$2.callBlockingMethod(MasterProtos.java)
 at org.apache.hadoop.hbase.ipc.RpcServer.call(RpcServer.java:387)
 at org.apache.hadoop.hbase.ipc.CallRunner.run(CallRunner.java:139)
 at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:369)
 at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:349)
```

None

**CDPD-56741: Improvement in log message when jwtauth not used**

The following exception is printed at startup only and it is not cluttering logs:

```
2023-05-30 06:18:40,127 ERROR org.apache.ranger.rms.security.RMSJwtAuthFilter:
 quasar-pibgzl-1.quasar-pibgzl.root.hwx.site-startStop-1]:
 Failed to initialize Ranger RMS JWT Auth Filter.
 java.lang.Exception: RangerJwtAuthHandler:
 Mandatory configs ('jwks.provider-url' & 'jwt.public-key')
 are missing, must provide atleast one.
 at org.apache.ranger.authz.handler.jwt.RangerJwtAuthHandler.initialize(RangerJwtAuthHandler.java:84)
 ~[ranger-authn-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
 at org.apache.ranger.rms.security.RMSJwtAuthFilter.initialize(RMSJwtAuthFilter.java:77)
 ~[ranger-rms-common-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
 at jdk.internal.reflect.NativeMethodAccessorImpl.invoke0(Native Method) ~[?:?]
 at jdk.internal.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl
```

This exception is logged, since mandatory JWT auth filter configs are not provided (expected for some environments like y-cloud) Even though JWT auth filter is failed to initialised, it falls back to kerberos auth, hence no impact from auth perspective

NA

**CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup**

This is a warning log printed in catalina.out file when Ranger RMS server is initialised. Following exception is observed only in Oracle 19 setup. FileNotFoundException: /usr/share/java/oraclepki.jar

NA

**CDPD-55107: Not able to search using multiple user filter in access audit tab**

If you were using multiple user search filters in Audit > Access Tab on Ranger Admin UI, after upgrading to CDP-7.1.9 that would not be supported. You can continue to search user with single search filter.

None

**CDPD-48975: Ranger KMS KTS to KMS DB migration : keys with the same name but different case are not migrated**

KMS keys are not case sensitive

No work arounds. Such keys combination are very rare and migration doc was updated to check such keys before starting the migration.

**CDPD-42598: Kafka policy creation allowed with incorrect permissions.**

When creating a Kafka policy from the UI, the permissions "Idempotent write" and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource

None.

**CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload:**

```
{"resourceName": "classification", "userInput": "", "resources": {"classification": []}}
```

expectation is to return all the classifications . But the response has only "classification" Happens similarly for entity-label , entity-business-metadata.

None.

**CDPD-40734: User allowed to insert data into a hive table when there is a deny policy on a table column.**

A user is allowed to enter data into a table even if there is a deny policy present on one of the table columns.

Test scenario details:

```
Policy setup :-
policy 1 :- all access policy for hrt_qa, hive and impala
users
resources - database - * , table - *, column - *
users : hrt_qa, hive, impala
access - all access allowed
policy 2 :- policy on test_1.table_1 for hrt_5
users : hrt_5
resources : database - test_1, table - table_1, column - *
access :- all access allowed
policy 3 :- deny policy on test_1.table_1.c0 for hrt_5
users : hrt_5
```

```
c0 resources : database - test_1, table - table_1, column -
 access - all access denied
 data setup :-
 database - test_1
 table - table_1(c0 int, c1 int)
```

The user is able to insert data into the table.

None.

### Known issues in Ranger KMS

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround.

#### CDPD-70115: Ranger KMS with Oracle DB not supported for Navigator Encrypt

Navigator Encrypt deposit registration is failing with Ranger KMS DB with Oracle DB setup with the following error:

```
java.sql.SQLException: ORA-02289: sequence does not
exist Error Code: 2289
```

None.

### Known Issues in Spark

Learn about the known issues in Spark the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7

The following known issues were identified in this release:

#### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7

##### CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the RuntimeWarning: Failed to add file message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

##### CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

**CDPD-67517: Spark3 tests are failing if /tmp is mounted as NOEXEC.**

Map the tmpdir to a writable path in spark3-conf/spark-defaults.conf using the following steps:

1. Go to the Cloudera Manager.
2. Go to the Spark service.
3. Click the Configuration tab.
4. Select Scope Gateway .
5. Select Category Advanced .
6. Locate the Spark Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-defaults.conf\_client\_config\_safety\_valve property.
7. Map the tmpdir to a writable path:

```
spark.driver.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
spark.executor.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
```

8. Enter a Reason for change, and then click Save Changes to commit the changes.
9. Deploy the client configuration.

**CDPD-23817: In the upgraded Cluster, the permission of /tmp/spark is restricted due to the HDP configuration hive.exec.scratchdir=/tmp/spark.**

If you are using the /tmp/spark directory in the CDP cluster, you must provide the required additional Policies/ACL permissions.

**CDPD-22670 and CDPD-23103: There are two configurations in Spark, "Atlas dependency" and "spark\_lineage\_enabled", which are conflicted. The issue is when Atlas dependency is turned off but spark\_lineage\_enabled is turned on.**

Run Spark application, Spark will log some error message and cannot continue. That can be restored by correcting the configurations and restarting Spark component with distributing client configurations.

**CDPD-23007: Mismatch in the Spark Default DB Location. In HDP 3.1.5, hive\_db entities have one attribute - 'location' which is configured to the '/managed' path. In fresh install of CDP 7.1.7, hive\_db entities now have 2 attributes 'location' configured to '/external' path and 'managedLocation' configured to '/managed' path. In the AM2CM migration (HDP 3.1.5 -> CDP 7.1.7), the 'location' attribute from hive\_db entities in HDP 3.1.5 comes unaltered to CDP 7.1.7 and hence maps to '/managed' path.**

This issue arises only if you are upgrading from HDP 3.1.5 to CDP 7.1.7. If you are performing a fresh install of CDP 7.1.7, you can ignore this issue.

None

**CDPD-217: The Apache Spark connector is not supported**

The old *Apache Spark - Apache HBase Connector* (shc) is not supported in CDP releases.

Use the new HBase-Spark connector shipped in CDP release.

**CDPD-3038: Launching pyspark displays several HiveConf warning messages**

When pyspark starts, several Hive configuration warning messages are displayed, similar to the following:

```
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of name hive.meta
store.runworker.in does not exist
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of
name hive.masking.algo does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of
name hive.metastore.runworker.in does not exist
```

```
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of na
me hive.masking.algo does not exist
```

These errors can be safely ignored.

### Known Issues in Schema Registry

Learn about the known issues in Schema Registry, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7

#### CDPD-40380: Authorization checking issue when Kerberos is disabled

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the `ranger.plugin.schema-registry.service.name` field.
3. Replace `GENERATED_RANGER_SERVICE_NAME` with the actual name of the service.
4. Restart the Schema Registry service.

#### CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

#### CDPD-59015: Schema Registry does not create new versions of schemas even if the schema is changed

Schema Registry uses a schema fingerprinting mechanism to differentiate between schemas. However, fingerprinting does not take into consideration the schema attributes of the field type. As a result, if you have two schemas where the only difference is that one has type attributes defined and the other does not, they will be considered identical by Schema Registry. For example, the following schemas are considered identical:

```
#Schema V1
{"type":"record","name":"schema_name","namespace":"ns","fields":
[{"name":"local_timestamp_micros_long","type":"long"}]}

#Schema V2
{"type":"record","name":"schema_name","namespace":"ns","fields":
[{"name":"local_timestamp_micros_long","type":{"type":"long",
"logicalType":"local-timestamp-micros"}]}]
```

Notice that the only difference is that in the second schema, the `local_timestamp_micros_long` field has a logical type specified. In cases like this, the new version of the schema is not created, the initial version is used. This is true even if the data that is being produced has a new schema version. The ID of the first schema version is used and is put in the serialized record. The new schema version is not created.

This issue is common when using change data capture (CDC) connectors like the Debezium connectors. This is because CDC connectors create schemas with the logical type decimal based on the column type in the database schema. For example:

```
{"type":"record","name":"schema_name","namespace":"ns","fields":
[{"name":"database_column","type":{"type":"bytes","logicalType":
"decimal","precision":64,"scale":0}}]}
```

If the database schema changes (for example, the column type), it is possible that only scale changes, which is a schema attribute.

```
{ "type": "record", "name": "schema_name", "namespace": "ns", "fields":
 [{ "name": "database_column", "type": { "type": "bytes", "logicalType":
 "decimal", "precision": 64, "scale": 1 } }] }
```

In this case, even though scale changed to 1, the first version of the schema is used where scale is 0. As a result, the data is consumed with the wrong scale.

Avoid using logical types or other attributes. Alternatively, ensure that there are no changes in the logical types or other attributes between schema versions.

#### **OPSAPS-68708: Schema Registry might fail to start if a load balancer address is specified in Ranger**

Schema Registry does not start if the address specified in the Load Balancer Address Ranger property does not end with a trailing slash (/).

Set the value of the RANGER\_REST\_URL Schema Registry environment variable to an address that includes a trailing slash.

1. In Cloudera Manager, select the Schema Registry service.
2. Go to Configuration.
3. Find the Schema Registry Server Environment Advanced Configuration Snippet (Safety Valve) property and add the following:

```
Key: RANGER_REST_URL
Value: [***RANGER REST API URL***]
```

Replace [\*\*\*RANGER REST API URL\*\*\*] with an address that can be used by Schema Registry to access Ranger. Ensure that the address ends with a trailing slash. For example: <http://ranger-1.cloudera.com:6182/>

4. Restart the Schema Registry service.

#### **OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade**

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

1. Access the Ranger Console (Ranger Admin web UI).
2. Click the cm\_atlas resource-based service.
3. Add the schemaregistry user to the all - \* policies.
4. Click Manage Service Edit Service .
5. Add the schemaregistry user to the default.policy.users property.

#### **OPSAPS-69317: Kafka Connect Rolling Restart Check fails if SSL Client authentication is required**

The rolling restart action does not work in Kafka Connect when the ssl.client.auth option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set ssl.client.auth to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

#### **Known Issues in Streams Replication Manager**

Learn about the known issues in Streams Replication Manager, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.



### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7

#### **CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic**

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

#### **CDPD-11079: Blacklisted topics appear in the list of replicated topics**

If a topic was originally replicated but was later excluded for replication, 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, this topic will be included in the response. Additionally, the excluded topic will also be visible in the SMM UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

#### **CDPD-30275: SRM may automatically re-create deleted topics on target clusters**

If `auto.create.topics.enable` is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with `srm-control`. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1]
```

2. Wait until SRM is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

### 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 in Cloudera Runtime 7.1.9 SP1 CHF 7.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7

#### **OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager**

Cloudera Manager does not display a Log Files menu for SMM UI role (and SMM UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by SMM UI is not supported by Cloudera Manager.

View the SMM UI logs on the host.

#### **OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners**

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL\_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in SMM safety valve in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server  
Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Add the  
following value for bootstrap servers>Save Changes > Restart SMM.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

#### **CDPD-70313: KNOX does not send Authentication header on FIPS configuration**

KNOX does not send neither the Authentication header nor the `hadoop.auth` cookie. Because of this, the SMM UI returns an HTTP 401 response and sets the `www-authenticate: Negotiate` header. After this, KNOX still does not send the Authentication header. This results in the SMM UI being inaccessible through Knox.

You can access the SMM UI directly in Cloudera Manager at `Clusters SMM Streams Messaging Manager Web UI` and log in using the Kerberos username and password.

#### **CDPD-82560: Streams Messaging Manager known unprotected endpoint**

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**

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**

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.

#### **Known Issues in Sqoop**

Learn about the known issues in Sqoop, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7**

There are no new known issues identified in this release.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7**

##### **CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version**

If you have upgraded to CDP PvC Base 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to CDP PvC Base 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see [Troubleshooting Apache Sqoop issues](#).

##### **CDPD-44431: Using direct mode causes problems**

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the `--direct` option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the `sqoop.enable.deprecated.direct` property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through `-Dsqoop.enable.deprecated.direct=true`.

#### **CDPD-3089: Avro, S3, and HCat do not work together properly**

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

#### **Parquet columns inadvertently renamed**

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.

#### **PARQUET-99: Importing Parquet files might cause out-of-memory (OOM) errors**

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.

#### **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 in Cloudera Runtime 7.1.9 SP1 CHF 7.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7**

There are no new known issues identified in this release.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7**

##### **COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."**

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

None.

##### **COMPX-4644: Queue capacity rounding problem when configuration is initially set via YARN**

When setting the capacity scheduler configuration through the YARN/Cloudera Manager configuration, there may be capacity values that use multiple decimal places. This results in rounding/floating point precision discrepancies in the UI when trying to validate that all sibling capacities equal 100%. The UI looks like all the numbers add up to 100, but the validation still displays an error and does not allow to save the capacities. It is also observed that the capacity is being calculated as, for example, 99.9999999991 in the backend.

- Create queues within the UI, or
- Ensure that capacities configured through the Capacity Scheduler safety valve do not have more than one decimal place.

#### **20202 Database migration after enabling opt-in migration**

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to CDP 7.1.9, forcing migration from H2 to PostgreSQL database.

- Upgrade to CDP 7.1.9 CHF2, moving back to H2 database.
- Upgrade to CDP 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version CDP 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

#### **CDPD-56559: MapReduce jobs can intermittently fail during a rolling upgrade.**

During a rolling upgrade between CDP versions 7.1.8 and 7.1.9, MapReduce jobs may fail with message, `RuntimeException: native snappy library not available`. Although the native Snappy compression library is not loaded, a checkmark displays to indicate that the Snappy compression library is loading for NodeManagers that are pending upgrades. This causes the MapReduce jobs that are associated with the NodeManagers to fail. After the upgrade, the jobs work as expected. This issue only impacts rolling upgrades from before CDP 7.1.9 to a higher version.

None.

#### **COMPX-12021 Queue Manager configurations on Scheduler Configuration page are not working**

When setting the following properties on the YARN Queue Manager UI, the properties are set in the `capacity-scheduler.xml` file which does not have any effect on YARN. The properties need to be set in the `yarn-site.xml` file, which does not happen when you set them through YARN Queue Manager.

- "Maximum Application Priority" – `"yarn.cluster.max-application-priority"`
- "Enable Monitoring Policies" – `"yarn.resourcemanager.scheduler.monitor.enable"`
- "Monitoring Policies" – `"yarn.resourcemanager.scheduler.monitor.policies"`



**Note:** You can also set this property on the YARN configuration page in Cloudera Manager as "Capacity Scheduler Preemption".

- "Preemption: Observe Only" – `"yarn.resourcemanager.monitor.capacity.preemption.observe_only"`
- "Preemption: Monitoring Interval (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.monitoring_interval"`
- "Preemption: Maximum Wait Before Kill (ms)" – `"yarn.resourcemanager.monitor.capacity.preemption.max_wait_before_kill"`
- "Preemption: Total Resources Per Round" – `"yarn.resourcemanager.monitor.capacity.preemption.total_preemption_per_round"`
- "Preemption: Over Capacity Tolerance" – `"yarn.resourcemanager.monitor.capacity.preemption.max_ignored_over_capacity"`
- "Preemption: Maximum Termination Factor" – `"yarn.resourcemanager.monitor.capacity.preemption.natural_termination_factor"`
- "Enable Intra Queue Preemption" – `"yarn.resourcemanager.monitor.capacity.preemption.intra-queue-preemption.enabled"`

1. In Cloudera Manager, select the YARN service.
2. Click the Configuration tab.
3. Search for `yarn-site.xml`.
4. Under YARN Service Advanced Configuration Snippet (Safety Valve) for `yarn-site.xml`, add the corresponding parameter and value you need.
5. Click Save Changes.
6. Restart the YARN services.

#### **COMPX-13177: QueueManager webapp requests fail with 'HTTP ERROR 400**

**`java.net.ConnectException: Unsupported ciphersuite TLS_EDH_RSA_WITH_3DES_EDE_CBC_SHA'`**

Products:

- Cloudera Manager for CDP Private
- Cloud Base Cloudera Manager for CDP Public Cloud

Context:

- Centos 7.8 and Redhat 7.8 operating systems, when FIPS support is enabled.

Problem:

- When attempting to display the Yarn Queue Manager interface, Cloudera Manager displays an error: "HTTP ERROR 400 java.net.ConnectException: Unsupported cipher suite 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.

### **Third-party applications do not launch if MapReduce framework path is not included in the client configuration**

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the `mapreduce.application.framework.path` property is set to the appropriate value, but third-party applications with their own configurations does not launch.

Set the `mapreduce.application.framework.path` property to the appropriate configuration for third-party applications.

### **JobHistory URL mismatch after server relocation**

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated the `mapred-site.xml` file that references the correct JobHistory Server.

### **YARN cannot start if Kerberos principal name is changed**

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN does not start. In such cases, the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the `znode` and restart the YARN service.
- Use the `reset ZK ACLs` command. This also sets the `znodes` below `/rmstore/ZKRMStateRoot` to `world:anyone:cdawa` which is less secure.

### **Queue Manager does not open on using a custom user with a default Kerberos principal**

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager `process_username` property matches the YARN `process_username` property.

### **Known Issues in Zeppelin**

Learn about the known issues in Zeppelin the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7**

The following known issues were identified in this release:

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7

#### 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.

#### BUG-125263: Zeppelin service move fails on clusters upgraded from HDP3.1.5

Resolve the circular symlink issue on the Zeppelin node by linking the conf directory to a new directory under /etc/zeppelin:

- # mkdir -p /etc/zeppelin/<version>/0
- # rm /usr/hdp/<version>/zeppelin/conf
- # ln -s /etc/zeppelin/<version>/0 /usr/hdp/<version>/zeppelin/conf

Where version is the HDP version. For example, 7.1.x-yyy. Restart the Zeppelin server in Ambari.

#### CDPD-3090: Due to a configuration typo, functionality involving notebook repositories does not work

Due to a missing closing brace, access to the notebook repositories API is blocked by default.

From the CDP Management Console, go to Cloudera Manager for the cluster running Zeppelin. On the Zeppelin configuration page (Zeppelin serviceConfiguration), enter shiro urls in the Search field, and then add the missing closing brace to the notebook-repositories URL, as follows:

```
/api/notebook-repositories/** = authc, roles[{{zeppelin_admin_group}}]
```

Click Save Changes, and restart the Zeppelin service.

#### CDPD-2406: Logout button does not work

Clicking the Logout button in the Zeppelin UI logs you out, but then immediately logs you back in using SSO.

Close the browser.

#### OPSAPS-59802: Zeppelin and Livy roles should be co-located on the same host.

When installing or upgrading to CDP Private Cloud Base, you must co-locate all Zeppelin and Livy roles on the same cluster host due to an issue with certificate generation.

### Known Issues in ZooKeeper

Learn about the known issues in ZooKeeper the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 7.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 7

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 7

There are no new known issues identified in the previous releases.

## Cumulative hotfix 6

Know more about the cumulative hotfix 6 for CDP 7.1.9 SP1. This cumulative hotfix was released on 07 Mar, 2025.

**Cloudera Runtime 7.1.9.1034 (Cumulative Hotfix 6) download URL:****Parcel Repository Location**

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1034/parcels/
```

**What's new in Cloudera Runtime 7.1.9 SP1 CHF 6**

Understand the functionalities and improvements to features of components in Cloudera Runtime 7.1.9 SP1 CHF 6.

**What's new in Knox****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](#).

**Fixed issues in 7.1.9 SP1 CHF 6**

Know more about the cumulative hotfix 6 for 7.1.9 SP1.



**Note:** Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p1034.63559640.

**CDPD-76868: Upgrade and unify commons-lang3 to version 3.13**

Upgraded the commons-lang3 dependency to version 3.13.0 throughout the stack.

**CDPD-75742: Upgrade commons-lang3 to 3.13.0**

The commons-text version 1.11.0 in Spark was calling Range.of(..) method of commons-lang3 that was not available in version 3.12.0 resulting in a Oozie-Spark job failure. This issue is now resolved and commons-lang3 is now upgraded to version 3.13.0 which contains the required method.

**CDPD-77188: Atlas can consume Iceberg CREATE TABLE statements from Cloudera Data Warehouse**

Incompatibilities between Cloudera Data Warehouse version 2024.0.18.2-4 with Cloudera Data Hub version 7.2.17 and Apache Atlas in Cloudera Runtime 7.1.9 CHF6 are resolved. Add the following property to the Atlas configuration to resolve the issue: atlas.notification.consumer.preprocess.iceberg.entities=true

**CDPD-77856: 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 fallbacks 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 has been added.

**CDPD-77736: Livy for Spark 3 Web UI was missing in the Knox gateway homepage in 7.1.8**

This issue is resolved by adding a metadata section to Livy for Spark 3 service.xml file. Knox can now display the appropriate section in its web UI.

**CDPD-78066: Iceberg tables when created, does not show under hive\_db entity**

After the update, the UI will show Iceberg tables if they have been created under the same hive\_db. Iceberg table can be seen in the **Tables** tab of a Hive\_db in basic search. Both hive\_table and iceberg\_table typenames will be visible.

**CDPD-78193: CSV injection vulnerability during CSV and Excel file export**

When policies were created with the special characters mentioned in a document, there were some vulnerabilities.

The issue is fixed now. Checks are now added to ensure whenever such characters are present, a space is added after it.

**CDPD-77092: 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.

**CDPD-79591: Impala base folder identification aligns with Hive**

Impala used a different approach to identify valid base folders for reading Hive ACID tables, leading to inconsistencies. Hive allowed reading a base folder even if an open writeId existed before a newer base writeId, but Impala did not. This discrepancy caused read issues, particularly when using INSERT OVERWRITE or TRUNCATE operations.

Impala's logic now treats base directories created by INSERT OVERWRITE and TRUNCATE the same way as Hive, ensuring consistent read behavior.

Apache Jira: [IMPALA-13759](#)

**CDPD-79456: 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-78850: Impala event processor now ignores CREATE\_TABLE events**

Previously, if a database was invisible to the Impala cluster due to authorization restrictions, CREATE\_TABLE events in such databases caused the event processor to enter an ERROR state.

Impala now ignores CREATE\_TABLE events when the database is not found, preventing errors in the event processor.

Apache Jira: [IMPALA-11735](#)

**CDPD-78836: vulnerable version of jquery.datatables is in use**

Upgrades to the latest supported version of jquery.datatables to address security issues.

**CDPD-79088: Impala now reloads metadata when table properties change in ALTER\_TABLE events**

After upgrading to Cloudera Data Platform 7.1.9 SP1 CHF 1, Impala automatically detected newly added partitions but did not refresh appended data in existing partitions unless a manual REFRESH or INVALIDATE METADATA command was run.

Impala now reloads metadata for ALTER\_TABLE events when table properties change, ensuring appended data is detected automatically.

Apache Jira: [IMPALA-13403](#)

**CDPD-79108: KerberosBasicAuthenticationHandler cannot handle colons in passwords**



KerberosBasicAuthenticationHandler can now handle colons in passwords. This handler is used by Schema Registry and Streams Messaging Manager.

**CDPD-72782: Ozone write does not work when http proxy is set for the JVM**

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**

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](#)

**CDPD-78069: Oozie action configuration's Java options are not applied due to CDPD-60551**

Java options configured in workflow xml's action configurations such as yarn.app.mapreduce.am.command-opts and mapreduce.map.java.opts are not applied to the Oozie Launcher AM's JVM. This issue is fixed now.

**COMPX-18589: YARN ResourceManager raised an exception during comparison of queues**

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-65645: Cache sort results in ContainerBalancerSelectionCriteria**

The sorting of all the containers was time consuming. The sort result is now cached and the performance of Container Balancer in large clusters is now improved.

Apache Jira: [HDDS-10160](#)

**CDPD-77399: HBase fails to register the servlet metrics and throws ClassNotFoundException: org.apache.hadoop.metrics.MetricsServlet**

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](#)

**CDPD-75089: Restrict trusted packages in ReflectData and SpecificData**

Schema parsing in the Java SDK of Apache Avro had an issue that could allow 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](#)

**CDPD-78468: Full outer join returns incomplete results**

Full outer join results were missing rows due to incorrect handling of filtered arrays across groups in the reducer.

The issue was addressed by ensuring that filtered arrays are correctly maintained per group, preventing data from one group from affecting another.

Apache Jira: [HIVE-23435](#)

**CDPD-77350: NullPointerException when selecting from a bucketed table with decimal column type**

Selecting data from a table with bucket column type as decimal and applying filter condition on the same bucket column throwing NullPointerException.

The issue was fixed by correctly handling bucket pruning for decimal columns, preventing the NullPointerException.

Apache Jira:[HIVE-28076](#)

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

- [CVE-2023-44483](#) - Xmlsec

### Known issues in 7.1.9 SP1 CHF 6

The following section lists the known issues and limitations, the areas of impact, and workaround for each Runtime component.

#### CDPD-78656: Health test for Knox fails if the `gateway.client.auth.needed = true` is set

The health test for Knox Gateway fails if the `gateway.client.auth.needed` parameter is set to true.

To avoid the startup failure, use `gateway.client.auth.wanted = true`.

#### CDPD-77911: Missing Log4j Redactor dependency

The missing `org.cloudera.logredactor` dependency can result impossible Log4j errors during startup. Additionally, it can lead to unwanted data leak because of incorrect log redaction.

#### CDPD-80921: Without permission for one glossary, `/glossary` call throws exception

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-79160: NPE while deleting BusinessMetadata

If business metadata is created without adding any applicable types, a `NullPointerException` is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

### Known Issues in Atlas

Learn about the known issues in Atlas, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6

#### CDPD-77054: Exporting shell entities can cause the import to fail

For Apache Hive tables created with Apache Spark, there may be shell entities created in Apache Atlas. If these entities are in the export zip file before they are resolved, the import will fail.

Workaround: None

#### CDPD-70450: Impala SQL queries that include the “WITH” clause should populate lineage in Atlas

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-6565: Issue with ddlQueries in Atlas for Hive/Impala tables created in default/non-default db using DAS/HUE respectively

ddlQueries are not created for Impala origin table. ddlQueries created for Impala/Hive CTAS tables have difference in names as compared to same operation for Hive table. The Impala table has the query in the ddlQueries entity name.

None

**CDPD-55301: ddlQueries and ALTERNATIVE\_\* lineage missing for Spark tables created through spark3-shell**

The ddlQueries and ALTERNATIVE\_\* lineage missing for Spark tables created using spark3-shell.

None

**CDPD-56085: 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**

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\*.

None

**CDPD-58554: Discard audits of specific classification, label, and business metadata**

Support to control audits for specific classification, label, and business metadata is not present in the Custom Audit Filters feature.

None

**CDPD-40346: The ddlQueries and ALTERNATIVE\_ADDCOLS lineage missing for Impala tables.**

The ALTERNATIVE\_ADDCOLS lineage has some issue when an Impala table is altered and the corresponding lineage is not created.

**CDPD-67112: Import transforms do not work as expected when replacing a string which already has ":"**

The character ":" is not supported in path replacements. The import succeeds but location remains unchanged. The character ":" must be avoided.

None

**CDPD-69150: Unable to add labels or user defined properties in Japanese**

Adding Japanese labels or user defined properties results in the error message: "Invalid label: ###, label should contain alphanumeric characters, \_ or -"

None

**CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties**

Entities, such as hdfs\_path, are not returned by the search when their indexable properties are searched using partial search terms made of Japanese or Chinese characters. Only exact matches return results when searching the indexed properties made of Japanese or Chinese characters.

None

**CDPD-68191: Suggestions do not return the correct results when searching multiple Chinese characters**

Free text search does not return results when searching Chinese phrases made of multiple characters. Partial searches return the correct results.

None.

**CDPD-71219: Regression : Suggestions don't work for single character words on indexable attributes**

When searching for entities whose name (entity names are indexable) is a single character, search results are returned but suggestions are not. When searching for entities whose description (entity descriptions are not indexable) is a single character, both search results and suggestions are returned.

None.

**CDPD-67450: Table name renaming operation is not updating or creating iceberg\_table entity**

Renaming an Iceberg table does not update the corresponding Atlas entity.

None.

**CDPD-67089: Export/Import: When a table with Ozone path is exported as "connected", only the Ozone key is exported.**

When table with Ozone path is exported as "connected", only Ozone key is exported. Other Ozone entities, such as Ozone volume, Ozone bucket are not exported.

None.

**CDPD-43772: Performance issues with Atlas service**

If there are lot of update operations and the compression type of column families of atlas\_janus table is SNAPPY, then the Kafka message processing might become slower.

- Consider setting compression type of column families of atlas\_janus table as GZ.

**OPSAPS-67783: During rolling upgrade one among two Atlas server failed to start but Cloudera Manager considered as success**

Cloudera Manager marks the Execute command Start on service Atlas-1 as a success even when Atlas service had failed to start successfully. In such cases, Atlas logs give the exact reason for Atlas start-up failure.

None

**CDPD-19358: "IsIndexable" and "isOptional" value of a typedef's attribute is modified post migration.**

In HDP-265, falcon\_feed\_creation has "stored-in" attribute which has IsIndexable value set to True.

Post migration, "stored-in" is moved to relationshipAttributeDefs and has IsIndexable value set to False.

Similarly, in HDP-265 the table attribute of hive\_storage\_desc isOptional is set to true. Post migration, isOptional set to False.

None.

**CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command**

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**

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**

In converting content from Navigator to Atlas, the create time for Hive tables is not moved to Atlas.

None

**CDPD-10574: Suggestion order doesn't match search weights**

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**

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) does not 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 the metadata collection from HMS. For example, if an ALTER TABLE operation in Spark that is changing a table name and is reporting 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 "@" does not 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, and use 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 does 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 logout 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 CDP.

None.

#### **CDPD-1664: Guest users are redirected incorrectly**

Authenticated users logging in to Atlas are redirected to the CDP 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: The 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, to add metadata tags to relationships that you wanted to make sure were unique in the system, you can design the relationship attribute with the property IsUnique equal true. However, in this release, the IsUnique attribute is not enforced.

None.

#### **DOCS-13759: Tag Propagation stops after a certain depth while the lineage is being extended**

When a tag is added to an entity at timestamp T1, the entities along the lineage to which the tag must be propagated is calculated at T1. Before tag propagation completes, if the lineage is extended, tag does not propagate to the entities in the extended lineage.

#### **CDPD-41142: When a Kafka console consumer group is run, more than one update audits are seen**

After running the console consumer with a consumer group, verify the consumer group entity created, along with the metrics and notifications for the consumer group and topic. The expected result can be: one ENTITY\_CREATE audit and one ENTITY\_UPDATE audit. But more than one ENTITY\_UPDATE audits are seen.

#### **CDPD-40165: Two audits are created for SPARK CTAS table**

When following Spark queries are fired:

```
spark.sql("create table table1(id int)")
spark.sql("create table table2 as select * from table1")
```

HMS sends "ENTITY\_CREATE" and "ENTITY\_FULL\_UPDATE\_V2".

The extra ENTITY\_FULL\_UPDATE\_V2 message received from HMS is sent as a part of ALTER TABLE ADDCOLS event from the HMS Hook side. This behaviour is observed only when the queries are run from Spark SQL and not when run the same queries are run from Beeline.

#### **CDPD-39197: Debug metrics returns empty data**

When debug metrics is enabled and some operations are performed, the response is empty

#### **CDPD-36495: Updating legacyAttribute from False to True resets the initially created relationshipAttributes values**

Creating types, entities, and to start, you must set the relationship with is\_legacy\_attribute value as **False**.

Later, update the value relationshipDef is\_legacy\_attribute to True.

For the entities that were created before updating the is\_legacy\_attribute to True, the relationship Attributes value is reset.

#### **CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities**

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

**CDPD-22744: Bulk entity DELETE API does not delete authorised entities**

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

**CDPD-29409: Hive import: Suggestion suggests entity which is deleted.**

Suggestions suggests tables of a database, which is a deleted entity.

**CDPD-25152: Tag propagation through deferred actions consumes additional time as compared to default flow**

The additional time might be due to the small overhead added to create / update task vertex and which is run in the background. This also depends on number of tasks queued to be executed.

**CDPD-42954: Zeppelin notebook fails after enabling Atlas-HDFS hook**

The Zeppelin notebooks are failing with errors after enabling Atlas-HDFS hook in the CDP cluster.

When the below properties are set for atlas-client.properties in Cloudera Manager:

- atlas.jaas.KafkaClient.option.keyTab
- atlas.jaas.KafkaClient.option.principal

Along with adding the properties in /etc/atlas/conf/atlas-application.properties, Cloudera Manager also adds these properties to atlas-application.properties for other services (like Spark).

Adding these properties interferes with the normal flow of the services (like Spark)

To enable HDFS lineage feature, instead of setting these properties through Cloudera Manager, users can manually add the properties directly in /etc/atlas/conf/atlas-application.properties

**CDPD-10576: Deleted Business Metadata attributes appear in Search Suggestions**

Atlas search suggestions continue to show Business Metadata attributes even if the attributes have been deleted.

None.

**CDPD-39427: [HDFS Lineage]: When the input is a directory in case of put/copyfromLocal/cp/mv commands, lineage is not created even though the script succeeds.**

When Source is a directory and target is a directory which is already present in Atlas, the command succeeds and inserts the data in the desired location, but lineage is not created.

**DOCS-13760: System Attributes search, \_\_classificationNames: Search with parent tag does not return entities associated to its children tags**

System attribute search with \_\_classificationNames = parent\_tag returns entities associated to parent\_tag only and not entities associated to its children tag.

Workaround: Instead of using system attribute, employ the basic search attribute "classification" which lists entities associated with inherited classifications.

**CDPD-35818: Basic search with tag filter provides approximateCount as -1 when there is no match and is 0 otherwise**

When the following search operations are performed:

- Faceted search with both tag filter and entity filter

The observed approximateCount is -1.

Here both entity filter and tag filter are present and when there is no match the response received is -1.

- Faceted search with only entity filter

Performing a basic search provides an approximate value of 0 when there is no match.

- Faceted search with only tag filter

Whenever there is tag filter in the query and there is no entity match, the approximateCount is -1 and if the tag filter is not available, the response approximateCount is 0

**CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities**

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

**CDPD-22744: Bulk entity DELETE API does not delete authorised entities**

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

**CDPD-76035: Resource lookup for Atlas service is failing**

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

**CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout**

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

**CDPD-79160: NPE while deleting BusinessMetadata**

If business metadata is created without adding any applicable types, a NullPointerException is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

**CDPD-82056: UI: when server response date fields as '0', UI shows as current time**

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](#)

**CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export**

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**

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](#)

**CDPD-66938: [Analyze] [Atlas] test\_time\_range tests fail**

When the Apache Atlas server is running on a node which has time zone other than UTC, there might be a time of day when the search results might differ if the relative CreateTime date range filters of TODAY, YESTERDAY, etc. are used. 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.

Use explicit date range filters instead of using relative date range filters, such as, TODAY, YESTERDAY.

**CDPD-84502: Advanced Search is not working properly in Atlas UI**



The Advance Search in Atlas does not return any results when selecting the Typedef option. The API is not triggered for each relation.

None

**CDPD-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails**

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

**CDPD-70321: Atlas Parallel import is failing with various errors**

During a parallel import-export activity with six iceberg table policies with exportOption as db1.\* for all six exports, all import fail after the exports.

**CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen**

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

**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**

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn**

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

**CDPD-79099: Unable to import hive tables using import script in Public Cloud**

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-88230: Bulk import operations currently do not construct lineage as expected**

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**

Bulk import operations do not create audits for the imported entities.

### Known issues in Navigator Encrypt

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround.

#### Failed to start navencrypt-mount.service

After installing NavEncrypt on SLES, if the command "systemctl status navencrypt-mount" fails with the error: "Failed to start navencrypt-mount.service: Unit navencrypt-mount.service failed to load: No such file or directory", the systemd files need to be reloaded.

Run the command:

```
(sudo) systemctl daemon-reload
```

#### Crash in Navigator Encrypt kernel module if --profile-file= option is used with ACLs

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 Key Trustee Server

Learn about the known issues in Key Trustee Server (KTS), the impact or changes to the functionality, and the workaround.

#### KTS setup on FIPS cluster gets stuck if KTS is in HA

Key Trustee Server setup gets stuck if KTS is in HA for Centos 7.9 in FIPS mode.

#### Create Date is not fetched from Key Trustee Server

During key migration from KTS to Ranger KMS, create date is not fetched from KTS.

### Known Issues in Ozone

Learn about the known issues in Ozone, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6

There are no new known issues identified in this release.

#### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6

##### 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: 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.

**The Prometheus binary is not available in CDP Private Cloud Base 7.1.9 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.

**After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and if Ozone is in the Non-HA environment, an exception message is observed during the finalization of the Ozone upgrade.**

During the finalization of the upgrade, a ClassNotFoundException message for org.cloudera.log4j.redactor.RedactorAppender class may be displayed. The error message can be ignored because the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

**CDPD-68951: In CDP Private Cloud Base 7.1.9 CHF2 version and lower, the command `ozone sh key list <bucket_path>` displays the `isFile` flag in a key's metadata as `false` even when the key is a file. This issue is rectified in CDP Private Cloud Base 7.1.9 CHF3. However, the pre-existing (pre-upgrade) key's metadata cannot be changed.**

None

**CDPD-67771: When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.**

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(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.

**CDPD-64398: In SCM with non-HA configuration, Secret key manager does not initialise. Hence, the startup of OM and DataNode fails because it cannot get a secret key. This key is used when security is enabled and used in Block token and container token verification while communicating between the Ozone client, OM, and Datanode.**

This issue does not occur in SCM with HA configuration.

You must force exit from the safe mode for SCM. This triggers the initialization of the Secret key manager. Below are the options:

- exit safe mode manually whenever SCM is started
- disable safe mode by setting `hdds.scm.safemode.enabled=false` in safety valve for SCM configuration

Impact of disabling safe mode and the purpose of safe mode:

- Write should not fail once the cluster is out of safe mode
- Read of existing data should not fail after the cluster is out of safe mode
- Unnecessary re-replication should be avoided during cluster restart

So disabling safe mode does not have any major impact on Ozone.

**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.

**CDPD-52412: keyManagerImpl#listStatus exceeds the maximum RPC length.**

To reduce the payload, you must increase the part size.

**OPSAPS-68159: If you did not deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Error when distributing parcel to host, Ignoring non-compliant parcel manifest error is displayed after Cloudera Manager is upgraded to 7.11.3.**

If you encounter the error, perform the following steps:

1. Deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.11.3.
2. Restart the cluster with a delay of 10 minutes.
3. Continue to upgrade CDH 7.1.8 to CDH 7.1.9.

**OPSAPS-68159: If you did not deactivate the Ozone parcel 718.2.x on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Ozone roles during the CDH 7.1.8 upgrade to CDH 7.1.9.**

If you encounter the error, perform the following steps:

1. Deactivate the Ozone parcel 718.2.x.
2. Restart the Ozone service.
3. Perform Finalize Upgrade for the Ozone service.

Step result: The Ozone roles will be displayed in green.

**CDPD-60989: The packaging version for Apache Ozone points to the 1.2.0 older version. This is a version string problem and not a packaging issue. The version of the Apache Ozone binary is closest to 1.4.0.**

None. This only affects the JAR versioning.

**CDPD-60366: Native library loader fails when system property `native.lib.tmp.dir` is not set. It fails because the library is copied to / instead of the cwd.**

Setting `native.lib.tmp.dir` to a certain path like `/tmp` should solve the issue. To set this system property, you must add `-Dnative.lib.tmp.dir=/tmp` to `ozone_java_opts` in Cloudera Manager configuration of the Ozone cluster.

**CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.**

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

**CDPD-60598: DataNode can consume up to the JVM heap size worth of direct memory buffers under specific failure scenarios. This can lead to a DataNode crash. This is due to Netty allocating and not freeing the direct memory buffers.**

Restarting DataNode will resolve this issue.

**CDPD-60466: The Ozone client is missing in the Cloudera Manager's Service Monitor node. This causes Ozone Canary's health check to fail.**

Ensure you install Ozone CLI on the same node as Cloudera Manager's Service Monitor.

**CDPD-60012: The Ozone SCM may be stuck in safe mode after upgrade to 7.1.9. This is due to a bug in how the SCM accounts for open storage containers.**

The workaround is to temporarily set the Ozone configuration `hdds.scm.safemode.threshold.pct` to a lower value like 0.90 and restart the SCM.

**CDPD-59679: In some scenarios, excess UNHEALTHY replicas of EC containers may not be removed from the cluster.**

None

**CDPD-50116: Topology-aware reads can provide better performance by routing applications to the closest replica of a file, if available, on the same physical node or same rack. However this feature is disabled by default in CDP 7.1.9.**

This feature can be enabled by setting `ozone.network.topology.aware.read` to true using Cloudera Manager safety valve.

**CDPD-57165: The DataNode disk usage thread may abort due to Ratis log tailing.**

Use `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory` to calculate disk usage. Configure in the `ozone-site.xml` property: `hdds.datanode.du.factory` value: `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory`

**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.

**CDPD-55513: Hive external table replication policy does not migrate Hive tables in HDFS storage to Ozone.**

To replicate data from HDFS to Ozone, see [Migrating your data from HDFS to Ozone](#). You can use HMS Mirror to replicate the metadata.

**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"`

#### **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_SHAREDLIB_DIR=/VAR/TMP`
3. Restart Ozone.

#### **CDPD-49137: Ozone Manager Kerberos token expires for SCM communication and OM does not log in again.**

Sometimes, OM's Kerberos token is not updated and it stops to communicate with SCM. When this occurs, writes start failing.

Restart OM or set the safety valve `hadoop.kerberos.keytab.login.autorenewal.enabled = true`

#### **CDPD-56684: Keys get deleted when you do not have permission on volume**

When a volume is deleted, it recursively deletes the buckets and keys inside it and only then deletes the volume. The volume delete ACL check is done only in the end, due to which you may end up deleting all the data inside the volume without having delete permission on the volume.



**Note:** There was no bucket/key permission available which allowed the user to delete them recursively.

#### **CDPD-50610: Large file uploads are slow with OPEN and stream data approach**

Hue file browser uses the append operation for large files. This API is not supported by Ozone in 7.1.9 and therefore large file uploads can be slow or timeout from the browser.

Use native Ozone client to upload large files instead of the Hue file browser.

#### **OPSAPS-64097: Ozone service restart failed at SCM**

Stopping SCM service using Cloudera Manager can sometimes timeout and need a retry. The Cloudera Manager API waits for 90 seconds which is not sufficient under certain circumstances.

Retry the shutdown using Cloudera Manager if the SCM still shows as running after a refresh of the service status.

#### **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-filesystem-hadoop3.jar`, it is likely due to the presence of the user `hdfs` on the node prior to CDP parcel

activation. As a result, the Cloudera Manager agent skips the initialization related to HDFS, leading to the non-creation of the `/var/lib/hadoop-hdfs` directory.

Create the directory `“/var/lib/hadoop-hdfs”` on all nodes followed by the deactivation and activation of the CDP parcel (deactivate and activate the Ozone parcel instead, in case Ozone parcel is used).

**CDPD-50447: When SCM High Availability is enabled, each of the SCM web UIs report the host of the web UI as the leader of HA, and the other two as followers. This gives wrong information**

Correct output is available by running the `ozone admin scm roles --service-id=<ID>` command.

**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](#).

**OPSAPS-62327: In an Ozone cluster without any gateway roles, Ozone is unable to deploy client configurations and displays the ConfigGenException error.**

You must add the Ozone gateway roles to the cluster.

**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-36389: The configurations `datanodes.involved.max.percentage.per.iteration` and `size.moved.max.per.iteration` are meant to limit the max number of DataNodes that is involved and maximum size that can move in an iteration. This bug causes balancer to stop an iteration when it is two DataNodes or 1 Container size (5GB) away from hitting these limits. However, these DataNodes can again be considered for balancing in the next iteration. This means the cluster is balanced after enough iterations, albeit a bit slowly. This bug is apparent in small clusters of around four DataNodes where the DataNodes can be either the source or target for a lot of moves but the iteration gets stopped when three DataNodes have been involved. It takes a higher number of iterations to eventually balance this cluster. While this is a performance issue, it does not prevent balancer from ultimately balancing the cluster.**

Increase the speed for balancing by decreasing the interval between each iteration using the configuration `balancing.iteration.interval`. The value of this configuration must be greater than `hdds.datanode.du.refresh.period`. `size.moved.max.per.iteration` can be increased to allow more data to move in one iteration.

**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-35141: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2> (state=08S01,code=40000) java.sql.SQLException: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2>. This issue does not occur if the source and target buckets are different in Hive queries. Currently, copying across the same bucket is only supported.**

Avoid different buckets in source and target path. Ensure that `hive.exec.stagingdir` is unset or set to the same bucket as the warehouse so that staging directories are created in the same bucket as the warehouse.

**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-41184: With Legacy buckets, FileSystem Optimised is not interoperating with the Ozone shell command. The cause is that the directory key entry in the DB KeyTable is stored as "dir1/" with trailing**



slash. But while performing the described operation, Ozone shell (o3://) is normalizing the given path and removing the trailing slash "/" from it. That results in KEY\_NOT\_FOUND exception.

There are three workarounds:

- Use FileSystem API to Delete the Directories rather than Shell-Command API.
- Use FSO buckets instead of Legacy Buckets. As in FSO, you can create Intermediate Directories and Delete Directories using the Ozone shell commands.
- Disable and set the configuration ozone.om.enable.filesystem.pathsflag to false in order to delete the directories. This is generally not a preferred workaround because the cluster must be restarted to pick up the new changes.

**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-12542: Mount of Ozone filesystem with the help of FUSE fails.**

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-42897: EC writes are failing with "No enough datanodes to choose" after EC replication config is set globally.**

EC writes start failing when large number of pipelines are created as a result of multiple EC configs with different chunk sizes used to write keys.

If standard EC configs (that is, rs-3-2-1024k) are used to write keys, number of pipelines created per DataNode is limited to five and this issue is not seen with standard EC configs.

The recommendation is not to create too many random chunk sizes. It is configurable because, users can decide based on their workload. But do not to have separate chunk sizes for each file.

**CDPD-41539: The "No such file or directory" message is displayed when EC file is read using older client.**

You must upgrade the client before trying to read the key: vol1/ecbuck1/1GB\_ec".

**CDPD-40560: Filesystem Operations using Hadoop S3a connector on a FILE\_SYSTEM\_OPTIMIZED bucket fails. org.apache.hadoop.ozone.om.exceptions.OMException: Unable to get file status: volume: s3v bucket: fso key: test/**

Do not run Hadoop S3a commands on a FILE\_SYSTEM\_OPTIMIZED bucket. Use OBJECT\_STORE bucket layouts.

**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.

**CDPD-43432: Ozone Service in fault state in DataNode - Long Running setup.**

Upgrade RocksDB to the latest version.

**OPSAPS-63999: In the newly installed cluster, the Finish upgrade option is clickable.**

None

**OPSAPS-64648: Failed to start Ozone node using Cloudera Manager if default log path /var/log/hadoop-ozone does not exist. If this path does not exist, any Ozone nodes (for example SCM or DataNode) restart will fail.**

Run the following command `sudo -u hdfs mkdir -p /var/log/hadoop-ozone` or replace hdfs with the user Ozone roles that are running.

**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.

## Technical Service Bulletins

**TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager**

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.

## Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

### Known issues in Ranger

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround.

#### UnsupportedClassVersionError

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.

#### CDPD-68806: The Revoke operation for users belonging to a group or role permission does not function as expected.

List command is listing all the tables even when the user permission is revoked. And also the command does not add any deny policy to Ranger for that specific user.

For example:

user\_1 is the part of group\_1

user\_2 is the part of role\_1

Create a Ranger policy for group\_1 and role\_1 to access all the resources in HBase.

Try to revoke the permission for user\_1 and user\_2 using HBase shell global revoke command:

```
revoke 'user_1'
```

```
revoke 'user_2'
```

But there is no deny policy added for that user. And access is still granted to user\_1 and user\_2.

This behavior is currently not supported in HBase shell. Must be handled manually using the Ranger policy change.

#### CDPD-68739: The revoke command does not work when using the HBase shell

While using the HBase shell, running the revoke command does not cancel the user permission. Users are able to perform actions even after running the revoke command.

For example:

When a user is granted permission to a particular table using the HBase grant command.

```
hbase:003:0> grant 'user_1', 'RWCA', 'some_abc_table'
```

And then execute the global revoke command for that user:

```
hbase:005:0> revoke 'user_1'
```

Permission for that user 'user\_1' for the table 'some\_abc\_table' not got revoked.

None.

#### CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

**CDPD-58860: After upgrading CDP-7.1.8 to CDP-7.1.9, cdp-proxy token missing from Knox Ranger policy.**

As part of [OPSAPS-67480](#) in 719 default ranger policy is added from cdp-proxy-token topology, so that after a new installation of CDP-7.1.9, the Knox Ranger policy includes cdp-proxy-token. However, upgrades do not add cdp-proxy-token to cm\_knox policies automatically.

Manually add cdp-proxy-token to the Knox policy, using Ranger Admin Web UI.

1. Log in to Cloudera Manager Ranger Admin Web UI, as a Ranger administrator.
2. On Ranger Admin Web UI Service Manager Resource Knox, click cm\_knox.
3. In Knox Policies, open the CDP Proxy UI, API and Token policy.
4. In Knox Topology\*, add cdp-proxy-token.
5. Click Save.
6. Restart Ranger.

**CDPD-58704: hadoop roll key / key delete command shows operation failed error when one KMS host is down, even when operation succeeds**

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered kms instances. If one KMS instance is stopped (not deleted), client gets 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.**

```
hbase:001:0> revoke 'hrt_11'
ERROR: org.apache.hadoop.hbase.coprocessor.CoprocessorException: HTTP 400 Error: processSecureRevokeRequest processing failed
 at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor.preRevoke(RangerAuthorizationCoprocesor.java:1309)
 at org.apache.ranger.authorization.hbase.RangerAuthorizationCoprocesor.preRevoke(RangerAuthorizationCoprocesor.java:1128)
 at org.apache.hadoop.hbase.master.MasterCoprocesorHost$162.call(MasterCoprocesorHost.java:1857)
 at org.apache.hadoop.hbase.master.MasterCoprocesorHost$162.call(MasterCoprocesorHost.java:1854)
 at org.apache.hadoop.hbase.coprocessor.CoprocessorHost$ObserverOperationWithoutResult.callObserver(CoprocessorHost.java:558)
 at org.apache.hadoop.hbase.coprocessor.CoprocessorHost.executeOperation(CoprocessorHost.java:631)
 at org.apache.hadoop.hbase.master.MasterCoprocesorHost.preRevoke(MasterCoprocesorHost.java:1854)
 at org.apache.hadoop.hbase.master.MasterRpcServices.revoke(MasterRpcServices.java:2740)
 at org.apache.hadoop.hbase.shaded.protobuf.generated.MasterProtos$MasterService$2.callBlockingMethod(MasterProtos.java)
 at org.apache.hadoop.hbase.ipc.RpcServer.call(RpcServer.java:387)
 at org.apache.hadoop.hbase.ipc.CallRunner.run(CallRunner.java:139)
 at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:369)
 at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExecutor.java:349)
```

None

#### CDPD-56741: Improvement in log message when jwtauth not used

The following exception is printed at startup only and it is not cluttering logs:

```
2023-05-30 06:18:40,127 ERROR org.apache.ranger.rms.security.RMS
JwtAuthFilter:
 quasar-pibgzl-1.quasar-pibgzl.root.hwx.site-startStop-1]:
 Failed to initialize Ranger RMS JWT Auth Filter.
 java.lang.Exception: RangerJwtAuthHandler:
 Mandatory configs ('jwks.provider-url' & 'jwt.public-key')
 are missing, must provide atleast one.
 at org.apache.ranger.authz.handler.jwt.RangerJwtAuthHan
 dler.initialize(RangerJwtAuthHandler.java:84)
 ~[ranger-authn-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
 at org.apache.ranger.rms.security.RMSJwtAuthFilter.init
 ialize(RMSJwtAuthFilter.java:77)
 ~[ranger-rms-common-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0
 -186]
 at jdk.internal.reflect.NativeMethodAccessorImpl.invoke
 0(Native Method) ~[?:?]
 at jdk.internal.reflect.NativeMethodAccessorImpl.invoke(N
 ativeMethodAccessorImp
```

This exception is logged, since mandatory JWT auth filter configs are not provided (expected for some environments like y-cloud) Even though JWT auth filter is failed to initialised, it falls back to kerberos auth, hence no impact from auth perspective

NA

#### CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup

This is a warning log printed in catalina.out file when Ranger RMS server is initialised. Following exception is observed only in Oracle 19 setup. FileNotFoundException: /usr/share/java/oraclepki.jar

NA

#### CDPD-55107: Not able to search using multiple user filter in access audit tab

If you were using multiple user search filters in Audit > Access Tab on Ranger Admin UI, after upgrading to CDP-7.1.9 that would not be supported. You can continue to search user with single search filter.

None

#### CDPD-48975: Ranger KMS KTS to KMS DB migration : keys with the same name but different case are not migrated

KMS keys are not case sensitive

No work arounds. Such keys combination are very rare and migration doc was updated to check such keys before starting the migration.

#### CDPD-42598: Kafka policy creation allowed with incorrect permissions.

When creating a Kafka policy from the UI, the permissions "Idempotent write" and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource

None.

#### CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload:

```
{"resourceName": "classification", "userInput": "", "resources": {"classification": []}}
```

expectation is to return all the classifications . But the response has only "classification" Happens similarly for entity-label , entity-business-metadata.

None.

#### **CDPD-40734: User allowed to insert data into a hive table when there is a deny policy on a table column.**

A user is allowed to enter data into a table even if there is a deny policy present on one of the table columns.

Test scenario details:

```
Policy setup :-
 policy 1 :- all access policy for hrt_qa, hive and impala
 users
 resources - database - * , table - * , column - *
 users : hrt_qa, hive, impala
 access - all access allowed
 policy 2 :- policy on test_1.table_1 for hrt_5
 users : hrt_5
 resources : database - test_1, table - table_1, column - *
 access :- all access allowed
 policy 3 :- deny policy on test_1.table_1.c0 for hrt_5
 users : hrt_5
 resources : database - test_1, table - table_1, column -
 c0
 access - all access denied
 data setup :-
 database - test_1
 table - table_1(c0 int, c1 int)
```

The user is able to insert data into the table.

None.

#### **Known issues in Ranger KMS**

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround.

#### **CDPD-70115: Ranger KMS with Oracle DB not supported for Navigator Encrypt**

Navigator Encrypt deposit registration is failing with Ranger KMS DB with Oracle DB setup with the following error:

```
java.sql.SQLException: ORA-02289: sequence does not
exist Error Code: 2289
```

None.

#### **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 in Cloudera Runtime 7.1.9 SP1 CHF 6

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6**

There are no new known issues identified in this release.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6**

**COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."**

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

None.

#### **COMPX-4644: Queue capacity rounding problem when configuration is initially set via YARN**

When setting the capacity scheduler configuration through the YARN/Cloudera Manager configuration, there may be capacity values that use multiple decimal places. This results in rounding/floating point precision discrepancies in the UI when trying to validate that all sibling capacities equal 100%. The UI looks like all the numbers add up to 100, but the validation still displays an error and does not allow to save the capacities. It is also observed that the capacity is being calculated as, for example, 99.9999999991 in the backend.

- Create queues within the UI, or
- Ensure that capacities configured through the Capacity Scheduler safety valve do not have more than one decimal place.

#### **20202 Database migration after enabling opt-in migration**

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to CDP 7.1.9, forcing migration from H2 to PostgreSQL database.
- Upgrade to CDP 7.1.9 CHF2, moving back to H2 database.
- Upgrade to CDP 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version CDP 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

#### **CDPD-56559: MapReduce jobs can intermittently fail during a rolling upgrade.**

During a rolling upgrade between CDP versions 7.1.8 and 7.1.9, MapReduce jobs may fail with message, RuntimeException: native snappy library not available. Although the native Snappy compression library is not loaded, a checkmark displays to indicate that the Snappy compression library is loading for NodeManagers that are pending upgrades. This causes the MapReduce jobs that are associated with the NodeManagers to fail. After the upgrade, the jobs work as expected. This issue only impacts rolling upgrades from before CDP 7.1.9 to a higher version.

None.

#### **COMPX-13177: QueueManager webapp requests fail with 'HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS\_EDH\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA'**

Products:

- Cloudera Manager for CDP Private
- Cloud Base Cloudera Manager for CDP Public Cloud

Context:

- Centos 7.8 and Redhat 7.8 operating systems, when FIPS support is enabled.

Problem:

- When attempting to display the Yarn Queue Manager interface, Cloudera Manager displays an error: "HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS\_EDH\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA".

1. Edit the file [/etc/default/cloudera-scm-server]

2. Around line 28, modify the line that starts with

```
#export CMF_OVERRIDE_TLS_CIPHERS=...
```

3. Remove the comment mark #.
4. Remove all ciphers with "3DES" in the name.
5. Save the file.
6. Restart the Cloudera Manager Server service.

### COMPX-12021 Queue Manager configurations on Scheduler Configuration page are not working

When setting the following properties on the YARN Queue Manager UI, the properties are set in the capacity-scheduler.xml file which does not have any effect on YARN. The properties need to be set in the yarn-site.xml file, which does not happen when you set them through YARN Queue Manager.

- "Maximum Application Priority" – "yarn.cluster.max-application-priority"
- "Enable Monitoring Policies" – "yarn.resourcemanager.scheduler.monitor.enable"
- "Monitoring Policies" – "yarn.resourcemanager.scheduler.monitor.policies"



**Note:** You can also set this property on the YARN configuration page in Cloudera Manager as "Capacity Scheduler Preemption".

- "Preemption: Observe Only" – "yarn.resourcemanager.monitor.capacity.preemption.observe\_only"
  - "Preemption: Monitoring Interval (ms)" – "yarn.resourcemanager.monitor.capacity.preemption.monitoring\_interval"
  - "Preemption: Maximum Wait Before Kill (ms)" – "yarn.resourcemanager.monitor.capacity.preemption.max\_wait\_before\_kill"
  - "Preemption: Total Resources Per Round" – "yarn.resourcemanager.monitor.capacity.preemption.total\_preemption\_per\_round"
  - "Preemption: Over Capacity Tolerance" – "yarn.resourcemanager.monitor.capacity.preemption.max\_ignored\_over\_capacity"
  - "Preemption: Maximum Termination Factor" – "yarn.resourcemanager.monitor.capacity.preemption.natural\_termination\_factor"
  - "Enable Intra Queue Preemption" – "yarn.resourcemanager.monitor.capacity.preemption.intra-queue-preemption.enabled"
1. In Cloudera Manager, select the YARN service.
  2. Click the Configuration tab.
  3. Search for yarn-site.xml.
  4. Under YARN Service Advanced Configuration Snippet (Safety Valve) for yarn-site.xml, add the corresponding parameter and value you need.
  5. Click Save Changes.
  6. Restart the YARN services.

### Third-party applications do not launch if MapReduce framework path is not included in the client configuration

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the mapreduce.application.framework.path property is set to the appropriate value, but third-party applications with their own configurations does not launch.

Set the mapreduce.application.framework.path property to the appropriate configuration for third-party applications.

### JobHistory URL mismatch after server relocation

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.



For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated the mapred-site.xml file that references the correct JobHistory Server.

#### **YARN cannot start if Kerberos principal name is changed**

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN does not start. In such cases, the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the znode and restart the YARN service.
- Use the reset ZK ACLs command. This also sets the znodes below /rmstore/ZKRMStateRoot to world:anyone:cdw which is less secure.

#### **Queue Manager does not open on using a custom user with a default Kerberos principal**

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager process\_username property matches the YARN process\_username property.

#### **Known Issues in HBase**

Learn about the known issues in HBase, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6**

##### **CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress**

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

##### **IntegrationTestReplication fails if replication does not finish before the verify phase begins**

During IntegrationTestReplication, if the verify phase starts before the replication phase finishes, the test will fail because the target cluster does not contain all of the data. If the HBase services in the target cluster does not have enough memory, long garbage-collection pauses might occur.

Workaround: Use the -t flag to set the timeout value before starting verification.

#### **HDFS encryption with HBase**

Cloudera has tested the performance impact of using HDFS encryption with HBase. The overall overhead of HDFS encryption on HBase performance is in the range of 3 to 4% for both read and update workloads. Scan performance has not been thoroughly tested.

Workaround: N/A

### AccessController postOperation problems in asynchronous operations

When security and Access Control are enabled, the following problems occur:

- If a Delete Table fails for a reason other than missing permissions, the access rights are removed but the table may still exist and may be used again.
- If `hbaseAdmin.modifyTable()` is used to delete column families, the rights are not removed from the Access Control List (ACL) table. The `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.

Workaround: N/A

Apache Issue: [HBASE-6992](#)

### Snappy compression with /tmp directory mounted with noexec option

Using the HBase client applications such as *hbase hfile* on the cluster with Snappy compression could result in `UnsatisfiedLinkError`.

Workaround: 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.

### HBase shutdown can lead to inconsistencies in META

Cloudera Manager uses an incorrect shutdown command. This prevents graceful shutdown of the HBase service and forces Cloudera Manager to kill the processes instead. It can lead to inconsistencies in Meta.

Workaround: Run the following command instead of shutting down the HBase service using Cloudera Manager.

```
hbase master stop --shutdownCluster
```

The command output must end with Closing master protocol: MasterService phrase. You can verify the command execution by checking the master logs. The log must contain Cluster shutdown requested of master=xxx and the closing of regions. Upon successful execution, the RegionServers start shutting down.



**Note:** The command does not stop the *REST Server* and the *Thrift Server* role instances. You can safely shut down them from Cloudera Manager later.

If you find any inconsistencies, please contact Cloudera Support.

### The custom JWT Principal claim is not usable for HBase

The `hbase.security.oauth.jwt.token.principal.claim` configuration property does not exist in the 7.1.9 SP1 version. The purpose of this property is to allow the use of a Subject or Principal claim different from the default sub claim.

Workaround: Use the standard sub claim in the JWT token.

### HBase procedures might hang because WriteAheadLog for ProcedureStore does not progress

HBase Master might hang and cause longer recovery when the WriteAheadLog is stuck while persisting to the ProcedureStore. HBase Master eventually aborts itself, however, the operations that rely on the ProcedureStore (for example, region movements and table creation) do not progress during this period.

Workaround: Restart the active HBase Master to speed up the recovery.

### CDPD-77399: HBase fails to register the servlet metrics and throws ClassNotFoundException: org.apache.hadoop.metrics.MetricsServlet

The MetricsServlet class is a Hadoop 2-based metric servlet unavailable in Hadoop 3 deployments.

Workaround: Ignore this WARN log message during HBase Master and RegionServer startup.

Apache Issue: [HBASE-28315](#)

### Known Issues in Hive

Learn about the known issues in Hive, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6

#### CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6

#### CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

#### CDPD-60770: Beeline Authentication Issue with Special Characters in Passwords

When LDAP is enabled, users cannot authenticate with Beeline if the password contains a special character. For example, the following string fails:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>;user=user@XXX;password='R3G#xpXyoy1M0Jb1'
```

Use the -p parameter to execute the Beeline command:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>; -n user@XXX -p 'R3G#xpXyoy1M0Jb1'
```

#### CDPD-43769: Inconsistent behavior with CHAR comparisons and string literals

In Hive, 'char' data to 'string' data comparison is considered as equal, but 'char' data to 'varchar' data comparison is considered as not equal if the data has trailing spaces. For example,

```
create table test(colchar char(10), colvarchar varchar(20), colstring string);
insert into test values ('a ', 'a ', 'a ');
select count(*) from test where colchar = colstring;
```

Output returns 1 because char and string data is considered as equal

```
select count(*) from test where colchar = colvarchar;
Output returns 0 because char and varchar data is considered as not equal
```

It is recommended that you use 'varchar or string' data type instead of 'char' data type for any character or string representation.

**CDPD-42726: Wrong results for agg queries from stats after running multi-insert query on managed table**

Multiple inserts at the same time into the same partition results in invalid stats. For example:

```
create table source(p int, key int, value string);
insert into source(p, key, value) values (101,42,'string42');

create table stats_part(key int,value string) partitioned by (p int);
from source
insert into stats_part select key, value, p
insert into stats_part select key, value, p;

select count(*) from stats_part;
-- In this scenario, StatsOptimizer helps serving this query be
cause the result should be rowNum of the partition p=101.
-- The result is 1, however, the result should be 2.
```

Rewrite the multi-insert statement by merging the branches inserting into the same partition to one branch.

**CDPD-57574: Query execution fails due to NullPointerException in DagUtils.setupQuickStart**

Hive queries may fail at runtime with a NullPointerException while executing the Tez graph at DagUtils.setupQuickStart method. A part of the stack trace from the error is shown below:

```
ERROR : Failed to execute tez graph.
java.lang.NullPointerException: null at org.apache.hadoop.hive.ql.exec.tez.DagUtils.setupQuickStart(DagUtils.java:1724)
```

The error usually occurs from an invalid plan that has cycles created by semi-join or map-join edges.

Disable some of the optimizations in order to remove the cycles from the plan. It is sufficient to set one of the following properties:

- set hive.tez.dynamic.partition.pruning=false;
- set hive.tez.dynamic.semijoin.reduction=false;
- set hive.auto.convert.join=false;

To avoid any undesired performance regressions, disable the optimization(s) only for specific queries and not globally (through Cloudera Manager) for the whole workload or cluster.

**CDPD-43107: SemanticException for INSERT INTO statement when Hive Cost-based Optimizer (CBO) is disabled**

If you are running the INSERT INTO query and hive.cbo.enable is set to "false", the query fails with a SemanticException. For example,

```
set hive.cbo.enable=false;

CREATE TABLE mytable (
id INT,
```

```
str STRING
);

INSERT INTO mytable (id, str) VALUES (1, 'a');
```

Output:  
 org.apache.hadoop.hive.ql.parse.SemanticException: 0:0 Expected 2 columns for insclause-0/default@mytable; select produces 1 columns. Error encountered near token ''a''

This issue occurs only when columns are specified for the target table in the query. INSERT INTO mytable values (1, 'a'); does not result in an exception.

Enable CBO and run the query.

#### CDPD-43957: HiveServer shuts down during replication due to high resource usage

During Hive replication (both bootstrap and incremental), you may notice that the HiveServer (HS2) shuts down periodically with the following error:

```
java.sql.SQLException: org.apache.hive.jdbc.ZooKeeperHiveClientException: Unable to read HiveServer2 configs from ZooKeeper
 at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.java:265)
 at org.apache.hive.jdbc.HiveDriver.connect(HiveDriver.java:107)
 at java.sql.DriverManager.getConnection(DriverManager.java:664)
 at java.sql.DriverManager.getConnection(DriverManager.java:247)
 at com.cloudera.enterprise.hive3qt.Hive3QueryTool$HiveOperation.execute(Hive3QueryTool.java:682)
 at com.cloudera.enterprise.hive3qt.Hive3QueryTool.main(Hive3QueryTool.java:935)
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Unable to read HiveServer2 configs from ZooKeeper
 at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:177)
 at org.apache.hive.jdbc.Utills.configureConnParamsFromZooKeeper(Utills.java:580)
 at org.apache.hive.jdbc.Utills.parseURL(Utills.java:391)
 at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.java:263)
 ... 5 more
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: Tried all existing HiveServer2 uris from ZooKeeper.
 at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.getServerHosts(ZooKeeperHiveClientHelper.java:132)
 at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureConnParams(ZooKeeperHiveClientHelper.java:172)
 ... 8 more
```

This issue occurs when you replicate at scale with an unbalanced cluster setup that has all the roles running on the Cloudera Manager host. As a result, Cloudera Manager ends up in a bottleneck situation because HS2 crashes, preventing further replication. The logs (/var/log/messages) indicate that there were a large number of Java processes running on the host, which exhausted the host's memory and triggered an Out Of Memory Killer process to stop HS2.

You can restart HS2 and reinitiate replication. However, the replication process may take a longer time to complete.

- Follow these recommendations when setting up your source and target clusters:
  - Ensure that no HS2, Hive metastore (HMS), DataNode, or NameNode roles are running on the Cloudera Manager host.
  - Ensure that you have multiple instances of HS2 and HMS roles on different nodes.
  - It is recommended that you have the HS2 and HMS roles on different nodes than the DataNodes or NameNodes.

These practices increase the possibility of Hive replication at scale completing successfully.

- Limit the scale of Hive replication by temporarily disabling replication policies.

#### **CDPD-40730: Parquet change can cause incompatibility**

Parquet files written by the parquet-mr library in CDP 7.1.9, where the schema contains a timestamp with no UTC conversion will not be compatible with older versions of Parquet readers. The effect is that the older versions will still consider these timestamps as they would require UTC conversions and will thus end up with a wrong result. You can encounter this problem only when you write Parquet-based tables using Hive, and tables have the non-default configuration `hive.parquet.write.int64.timestamp=true`.

None.

#### **CDPD-41274: HWC + Oozie issue: Could not open client transport with JDBC Uri**

Currently only Spark cluster mode is supported in the Oozie Spark Action with Hive Warehouse Connector (HWC).

Use Spark action in cluster mode.

```
<spark xmlns="uri:oozie:spark-action:1.0">
 ...
 <mode>cluster</mode>
 ...
</spark>
```

#### **CDPD-26556 After an upgrade, querying a CTAS table under certain conditions might throw an exception**

If you upgrade your Hive cluster from CDH 6 to CDP 7, create a CTAS table in the CDP cluster from a table you upgraded from CDH, you might see the following exception when you query the new table:

```
class org.apache.hadoop.io.IntWritable cannot be cast to class o
rg.apache.hadoop.hive.serde2.objectinspector.StandardUnionObject
Inspector$StandardUnion
```

This issue involves CDH-based tables having columns of complex types ARRAY, MAP, and STRUCT.

#### **CDPD-23506: OutOfMemoryError in LLAP**

Long running spark-shell applications can leave sessions in interactive Hiveserver2 until the Spark application finishes (user exists from spark-shell), causing memory pressure in case of a high number of queries in the same shell (1000+).

You must close spark-shell so that sessions are closed. Add the owner of the database or the tables as a user with read or read/write access to the tables directly.

#### **CDPD-23041: DROP TABLE on a table having an index does not work**

If you migrate a Hive table to CDP having an index, DROP TABLE does not drop the table. Hive no longer supports indexes ([HIVE-18448](#)). A foreign key constraint on the indexed table prevents dropping the table. Attempting to drop such a table results in the following error:

```
java.sql.BatchUpdateException: Cannot delete or update a parent
row: a foreign key constraint fails ("hive"."IDXS", CONSTRAINT "
```

```
IDXS_FK1" FOREIGN KEY ("ORIG_TBL_ID") REFERENCES "TBLS ("TBL_ID"
))
```

There are two workarounds:

- Drop the foreign key "IDXS\_FK1" on the "IDXS" table within the metastore. You can also manually drop indexes, but do not cascade any drops because the IDXS table includes references to "TBLS".
- Launch an older version of Hive, such as Hive 2.3 that includes IDXS in the DDL, and then drop the indexes as described in [Language Manual Indexing](#).

Apache Issue: [Hive-24815](#)

#### **CDPD-45370: Cartesian join fails with an empty table**

When the configuration `hive.tez.cartesian-product.enabled` is set to true, query fails if one of the tables involved in a Cartesian join has zero rows. This results in an `IllegalArgumentException` with the message `Number of items is 0. Should be positive`.

To address this issue, set `hive.tez.cartesian-product.enabled` to false before running the query.

Apache Jira: [\[HIVE-26582\]](#)

#### **CDPD-17766: Queries fail when using `spark.sql.hive.hiveserver2.jdbc.url.principal` in the JDBC URL to invoke Hive.**

Do not specify `spark.sql.hive.hiveserver2.jdbc.url.principal` in the JDBC URL to invoke Hive remotely.

Workaround: specify `principal=hive.server2.authentication.kerberos.principal` as shown in the following syntax:

```
jdbc:hive://<host>:<port>/<dbName>;principal=hive.server2.authentication.kerberos.principal;<otherSessionConfs>?<hiveConfs>#<hiveVars>
```

#### **CDPD-13636: Hive job fails with `OutOfMemory` exception in the Azure DE cluster**

Set the parameter `hive.optimize.sort.dynamic.partition.threshold=0`. Add this parameter in Cloudera Manager (Hive Service Advanced Configuration Snippet (Safety Valve) for `hive-site.xml`)

#### **CDPD-10848: HiveServer Web UI displays incorrect data**

If you enabled auto-TLS for TLS encryption, the HiveServer2 Web UI does not display the correct data in the following tables: Active Sessions, Open Queries, Last Max n Closed Queries

### **Technical Service Bulletins**

#### **TSB 2021-501: JOIN queries return wrong result for join keys with large size in Hive**

JOIN queries return wrong results when performing joins on large size keys (larger than 255 bytes). This happens when the fast hash table join algorithm is enabled, which is enabled by default.

#### **Knowledge article**

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2021-501: JOIN queries return wrong result for join keys with large size in Hive](#)

#### **TSB 2023-702: Potential wrong result for queries with date partition filter for clusters in GMT+ timezone**

In Cloudera Data Platform (CDP) Private Cloud Base 7.1.7 Service Pack (SP) 2 Cumulative Hotfix (CHF) 11, a fix was introduced in Hive Metastore (HMS) to address a parsing issue with date strings. This fix caused a regression in Hive clusters where the HMS time zone is set ahead of GMT for the following combination of tables and queries: a table that is partitioned on a DATE column and a SELECT query on that table containing a WHERE clause filter on the same DATE column. For such queries, during the partition pruning phase, the date string would be converted to a date

without timezone and compared with the partition value retrieved by HMS. This causes wrong results (0 rows) because the date values do not match.

The regression was identified in CDP Private Cloud Base 7.1.7 SP2 CHF14, but it exists in CHF11 through CHF16 as well as on certain versions of 7.1.8 and 7.1.9.

This issue does not affect clusters where the time zones are behind GMT. For example, if the time zone of the cluster is set USA/Los Angeles, which is 8 hours behind GMT, a date '2023-10-02' will remain as '2023-10-02' after converting to GMT (adding 8 hours). On the other hand, using Asia/Hong Kong time as an example, which is 8 hours ahead of GMT, the same date would become '2023-10-01' after converting to GMT (subtracting 8 hours), which leads to the wrong results.

### Upstream JIRA

[HIVE-27760](#)

### Knowledge article

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2023-702: Potential wrong result for queries with date partition filter for clusters in GMT+ timezone](#)

### Known Issues in Iceberg

Learn about the known issues in Iceberg, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6

There are no known issues identified in earlier releases.

### Known Issues in Impala

Learn about the known issues in Impala, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6

#### CDPD-49323: INSERT statement does not respect Ranger policies for HDFS

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm\_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXX XXXX

Apache Jira: [IMPALA-11871](#)

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6

#### CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka



Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

#### **CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers**

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the `--tls_ciphersuites` startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.

Ensure that the list in the `--tls_ciphersuites` startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

**Apache Jira:** [IMPALA-14625](#)

#### **CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.**

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

#### **CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work**

If you try to run `impala-shell` on RHEL 9 by setting the default python executable available in `PATH` to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, `IMPALA_PYTHON_EXECUTABLE=python3`.

#### **CDPD-42958: After upgrading the CDH 7.1.9 from CDH 6.x, under certain conditions you cannot insert data into a table**

Under the following conditions, after upgrading from CDH 6.x to CDH 7.1.9 you cannot insert data into a table from Impala:

- On CDH 6.x, you created a database with Impala in a user specified HDFS location.
- Using Hive, you then created a table in the database.

Under these conditions, the database and table are stored in the user-specified HDFS directory. After upgrading, the HDFS directory of the table is read-only for Impala. Consequently, from Impala you cannot insert new data into the table because Impala does not have write permission on the HDFS directory.

Workaround: To resolve this issue, use either one of the following workarounds:

- Using the Ranger Web UI, in the policy repository `cm_hdfs`, grant the user 'impala' write permission on the directory where the table resides.
- Enter the following command to grant write permission to user 'impala' on the HDFS directory where the table resides.

```
hdfs dfs -setfacl -m default:user:impala:rwx <HDFS directory>
```

#### **CDPD-80019: Impala fails to start due to large catalog-update topic**

Impala coordinators fail to start if the catalog-update statestore topic exceeds 2GB, hitting the Thrift 0.16.0 message size limit (`thrift_rpc_max_message_size`). This prevents the coordinator from processing the update.

- Enable local catalog mode: `--use_local_catalog=true`
- In Cloudera Manager, go to Impala Configuration search for "Catalog Server Command Line Argument Advanced Configuration Snippet (Safety Valve)" and click 
- Add property name and value:

```
Name : --catalog_topic_mode
Value: minimal
```

- Save changes.

Apache Jira: [IMPALA-12350](#)

### 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 HDF
S file hdfs://nameservice1/warehouse/tablespace/managed/hive/<da
tabase>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException:
File does not exist: /warehouse/tablespace/managed/hive/<data
base>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

### CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

You must use a single Knox node to access Impala UI.

### Impala api calls via kinox require configuration if the kinox customized kerberos principal name is a default service user name

To access impala api calls via kinox, if the kinox customized kerberos principal name is a default service user name, then configure "authorized\_proxy\_user\_config" by clicking Clusters->impala->configuration. Include the kinox customized kerberos principal name in the comma separated list of values <knox\_custom\_kerberos\_principal\_name>="\*" where <knox\_custom\_kerberos\_principal\_name> is the value of the Kerberos Principal in the Knox service. Select Clusters>Knox>Configuration and search for Kerberos Principal to display this value.

### CDPD-21828: Multiple permission assignment through grant is not working

None

### Problem configuring masking on tables using Ranger

The following Knowledge Base article describes the behavior when we configure masking on tables using Ranger. This configuration works for Hive, but breaks queries in some scenarios for Impala.

For a workaround, see the following Knowledge Base article: [ERROR: "AnalysisException: No matching function with signature: mask\(FLOAT\)" when Impala jobs fail with the following error with signature: mask\(FLOAT\)](#)

### IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm\_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions

happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXXXX"

#### **IMPALA-532: Impala should tolerate bad locale settings**

If the LC\_\* environment variables specify an unsupported locale, Impala does not start.

Add LC\_ALL="C" to the environment settings for both the Impala daemon and the Statestore daemon.

#### **Avro Scanner fails to parse some schemas**

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

#### **IMPALA-691: Process mem limit does not account for the JVM's memory usage**

Some memory allocated by the JVM used internally by Impala is not counted against the memory limit for the impalad daemon.

To monitor overall memory usage, use the top command, or add the memory figures in the Impala web UI /memz tab to JVM memory usage shown on the /metrics tab.

#### **IMPALA-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon**

If an Avro table has a schema definition with a trailing semicolon, Impala encounters an error when the table is queried.

Remove trailing semicolon from the Avro schema.

#### **IMPALA-1652: Incorrect results with basic predicate on CHAR typed column**

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the RPAD() function to blank-pad literals compared with CHAR columns to the expected length.

#### **IMPALA-1821: Casting scenarios with invalid/inconsistent results**

Using a CAST() function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

#### **IMPALA-2005: A failed CTAS does not drop the table if the insert fails**

If a CREATE TABLE AS SELECT operation successfully creates the target table but an error occurs while querying the source table or copying the data, the new table is left behind rather than being dropped.

Drop the new table manually after a failed CREATE TABLE AS SELECT

#### **IMPALA-2422: % escaping does not work correctly when occurs at the end in a LIKE clause**

If the final character in the RHS argument of a LIKE operator is an escaped \% character, it does not match a % final character of the LHS argument.

None

#### **IMPALA-2603: Crash: impala::Coordinator::ValidateCollectionSlots**

A query could encounter a serious error if includes multiple nested levels of INNER JOIN clauses involving subqueries.

None

#### **IMPALA-3094: Incorrect result due to constant evaluation in query with outer join**

An OUTER JOIN query could omit some expected result rows due to a constant such as FALSE in another join clause. For example:

```

explain SELECT 1 FROM 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 |
+---+

```

### IMPALA-3509: Breakpad minidumps can be very large when the thread count is high

The size of the breakpoint minidump files grows linearly with the number of threads. By default, each thread adds 8 KB to the minidump size. Minidump files could consume significant disk space when the daemons have a high number of threads.

Add `-minidump_size_limit_hint_kb=size` to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then 2 KB per thread after that.) The minidump file can still grow larger than the "hinted" size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

## IMPALA-4978: Impala requires FQDN from hostname command on Kerberized clusters

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the `gethostname()` system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a Kerberized cluster.

Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname`, only returns the short name, pass the command-line flag `##hostname=FULLY_QUALIFIED_DOMAIN_NAME` in the startup options of all Impala-related daemons.

## IMPALA-6671: Metadata operations block read-only operations on unrelated tables

Metadata operations that change the state of a table, like `COMPUTE STATS` or `ALTER RECOVER PARTITIONS`, may delay metadata propagation of unrelated unloaded tables triggered by statements like `DESCRIBE` or `SELECT` queries.

None

## IMPALA-7072: Impala does not support Heimdal Kerberos

None

**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.

**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.

**Known Issues in Knox**

Learn about the known issues in Knox, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6.

**Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6**

There are no new known issues identified in this release.

**Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6****CDPD-61527: On FIPS cluster, Knox-Impala connection is failing with SSL error code 5**

The Knox-Impala connection is failing with SSL error code 5 on FIPS clusters. This prevents accessing Impala through Knox.

**CDPD-78656: Health test for Knox fails if the gateway.client.auth.needed = true is set**

The health test for Knox Gateway fails if the gateway.client.auth.needed parameter is set to true. Environments using the "curl" call are impacted. The curl call from CM is not specifying any certificate (store) while Knox is configured to require one.

To avoid the startup failure, use gateway.client.auth.wanted = true.

**Known Issues in Livy**

Learn about the known issues in Livy the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6.

**Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6**

The following known issues were identified in this release:

No new known issues in this release.

**Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6**

No known issues identified in Livy.

**Known Issues in Hue**

Learn about the known issues in Hue, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6

**Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6**

There are no new known issues identified in this release.

**Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6****CDPD-58978: Batch query execution using Hue fails with Kerberos error**

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark

queries in a batch mode launches a shell script Oozie job from Hue and this is not supported on a secure cluster.

There is no workaround. You can submit the queries individually.

#### **OPSAPS-69659: Hue service fails on restart with "Unable to find psycopg2 2.5.4" error**

Hue service fails to restart and you see the following error: Unable to find psycopg2 2.5.4. This could be because you have installed Python in a non-default location and Hue is unable to locate the psycopg2 PostgreSQL database adapter.

You must specify the path where you have installed Python in the PYTHONPATH property in the Hue Advanced Configuration Snippet using Cloudera Manager.

1. Log in to Cloudera Manager as an Administrator.
2. Go to Clusters Hue Configurations and add the following key and value in the Hue Service Environment Advanced Configuration Snippet (Safety Valve) field:

Key: PYTHONPATH

Value: `[***PYTHON-PATH***]`

Replace `[***PYTHON-PATH***]` with the actual location where you have installed Python. For example, `/opt/cloudera/parcels/CDH/lib/hue/build/env/lib/python3.8/site-packages`

3. Click Save Changes.
4. Restart the Hue service.

#### **CDPD-59677: Unable to view Phoenix tables on the left assist in Hue**

On clusters secured with Knox, you may not be able to see Phoenix tables on the left assist that are present under the default database (that is, an empty("") database).

None.

#### **CDPD-58142: A query is not pre-populated in the Hue editor after clicking on the "Re Execute" button**

When you click Re Execute to rerun a query from the Job Browser Queries Query Details page, the query does not get populated on the Hue editor, as expected.

None.

#### **CDPD-54376: Clicking the home button on the File Browser page redirects to HDFS user directory**

When you are previewing a file on any supported filesystem, such as S3 or ABFS, and you click on the Home button, you are redirected to the HDFS user home directory instead of the user home directory on the said filesystem.

None.

#### **CDPD-41306: pip3.8 freeze command does not work and results into an error**

You see the `/usr/lib/hue/build/env/bin/python: No such file or directory` error when you run the following command:

```
build/env/bin/pip3.8 freeze
```

Run the freeze command as follows by specifying the paths of python3.8 and pip3.8:

```
/opt/cloudera/parcels/CDH/lib/hue/build/env/bin/python3.8 /opt/cloudera/parcels/CDH/lib/hue/build/env/bin/pip3.8 freeze
```

#### **CDPD-43293: Unable to import Impala table using Importer**

Creating Impala tables using the Hue Importer may fail.

If you have both Hive and Impala services installed on your cluster, then you can import the table using by selecting the Hive dialect from Tables Sources .

If only Impala service is installed on your cluster, then go to Cloudera Manager Clusters Hue Configurations and add the following line in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue\_safety\_valve.ini field:

```
[beeswax]
max_number_of_sessions=1
```

**DWX-8602: Unable to import a large CSV file from the local workstation**

You may see an error message while importing a CSV file into Hue from your workstation, stating that you cannot import files of size more than 200 KB.

Upload the file to S3 or ABFS and then import them into Hue using the Importer.

**CDPD-39330: Unable to use the pip command in CDP**

You may not be able to use the pip command in CDP 7.1.7 or higher and may see the following error when using pip in a command: “ImportError: cannot import name chardet”.

Follow the steps listed on [Unable to use pip command in CDP](#).

**CDPD-24294: Hue uses the unsafe-inline directive in its Content Security Policy (CSP) header**

Hue 4 web interface uses the unsafe-inline directive in its CSP header. As a result, the application server does not set the CSP header in its HTTP responses, and therefore does not benefit from the additional protection against potential cross-site scripting issues and other modern application vulnerabilities which a properly configured CSP may provide. This could lead to application vulnerability.

Cloudera recommends deploying additional security measures such as a firewall within the Hue server to control allowed connections, and SSO-based authentications mechanisms such as LDAP or SAML.

**OPSAPS-58927: Connection failed error when accessing the Search app (Solr) from Hue**

If you are using Solr with Hue to generate interactive dashboards and for indexing data, and if you have deployed two Solr services on your cluster and selected the second one as a dependency for Hue, then Cloudera Manager assigns the hostname of the first Solr service and the port number of the second Solr service generating an incorrect Solr URL in the search section of the hue.ini file. As a result, you may see a “Connection failed” error when you try to access the Search app from the Hue web UI.

1. Log into Cloudera Manager as an Administrator.
2. Go to Clusters Hue service Configuration and add the following lines in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue\_safety\_valve.ini field:

```
[search]
URL of the Solr Server
solr_url=http://[***HOSTNAME***]:[***PORT***]/solr/
```

For example:

```
solr_url=http://solr2:4567/solr/
```

3. Click Save Changes.
4. Restart the Hue service.

**CDPD-41136: Importing files from the local workstation is disabled by default**

Cloudera has disabled the functionality to import files from your local workstation into Hue because it may cause errors. You may not see the Local File option in the Type drop-down menu on the **Importer** page by default.

You can enable the functionality to import files from your local workstation by specifying the following parameter in the Hue Service Advanced Configuration Snippet (Safety Valve) for hue\_safety\_valve.ini field in Hue configurations in Cloudera Manager:

```
[indexer]
enable_direct_upload=true
```

### Invalid S3 URI error while accessing S3 bucket

The Hue Load Balancer merges the double slashes (//) in the S3 URI into a single slash (/) so that the URI prefix "/filebrowser/view=S3A:/" is changed to "/filebrowser/view=S3A:". This results in an error when you try to access the S3 buckets from the Hue File Browser through the port 8889.

The Hue web UI displays the following error: “Unknown error occurred”.

The Hue server logs record the “ValueError: Invalid S3 URI: S3A” error.

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.

You should be able to load the S3 browser from both 8888 and 8889 ports.

Alternatively, you can use the Hue server port 8888 instead of the load balancer port 8889 to resolve this issue.

### Error while rerunning Oozie workflow

You may see an error such as the following while rerunning an already executed and finished Oozie workflow through the Hue web interface: E0504: App directory [hdfs://cdh/user/hue/oozie/workspaces/hue-oozie-1571929263.84] does not exist.

To resolve this issue, add the following property in the Hue Load Balancer Advanced Configuration Snippet:

1. Sign in to Cloudera Manager as an administrator.
2. Go to Clusters Hue service Configurations Load Balancer and search for the Load Balancer Advanced Configuration Snippet (Safety Valve) for httpd.conf field.
3. Specify MergeSlashes OFF in the Load Balancer Advanced Configuration Snippet (Safety Valve) for httpd.conf field.
4. Click Save Changes.
5. Restart the Hue Load Balancer.

### CDPD-17956, CDPD-14167: Setting idle session timeout for Hue does not work when the cluster is secured using Knox SSO

If Hue is configured with desktop.auth.backend.KnoxSpnegoDjangoBackend as the Authentication Backend, then the automatic idle session logout that is set by configuring the idle\_session\_timeout property does not take effect. You may also see 404 error while accessing Hue from the Knox UI when the idle\_session\_timeout property is not set to -1.

None

### Query history displays "Result Expired" message

You see the "Result Expired" message under the Query History column on the **Queries** tab for queries which were run back to back. This is a known behaviour.



None.

#### **CDPD-43956: Manually replace UUID when importing Oozie workflows containing sub-workflows**

When importing Oozie workflows that contain sub-workflows, you must replace all the UUID entries with unique new entries. If you change the UUID of a sub-workflow, you must update that reference in the parent workflow to avoid circular dependencies.

#### **Hue limitation after upgrading from CDH to CDP Private Cloud Base**

The `hive.server2.parallel.ops.in.session` configuration property changes from `TRUE` to `FALSE` after upgrading from CDH to CDP Private Cloud Base. Current versions of Hue are compatible with this property change; however, if you still would like to use an earlier version of Hue that was not compatible with this property being `FALSE` and shared a single JDBC connection to issue queries concurrently, the connection will no longer work after upgrading.

#### **Hue UI is blank upon login after upgrading to CDP 7.1.7 from CDH 6**

If your cluster was secured using Knox, and if you have upgraded from CDH 6 to CDP 7.1.7, then you may see a blank Hue screen. This could happen because the `knox_proxyhosts` parameter is newly introduced in CDP, and it is possible that this parameter is not configured in Cloudera Manager under Hue configuration.

Specify the host on which you have installed Knox in the Hue Knox Proxy Hosts configuration as follows:

1. Log in to Cloudera Manager as an Administrator.
2. Obtain the host name of the Knox Gateway by going to Clusters Knox service Instances .
3. Go to Clusters Hue service Configuration and search for the Knox Proxy Hosts field.



**Note:** Cloudera Manager displays the following warning if the Knox Proxy Hosts field is empty when Knox Gateway is enabled on the CDP cluster: The parameter `knox_proxyhosts` cannot be empty. This can happen if there are no Knox Gateways. Please set the `knox_proxyhosts` to the list of hosts that have Knox Gateways.

4. Specify the Knox Gateway hostname in the Knox Proxy Hosts field.
5. Click Save Changes and restart the Hue service.

### **Unsupported features**

#### **CDPD-59595: Spark SQL does not work with all Livy servers that are configured for High Availability**

SparkSQL support in Hue with Livy servers in HA mode is not supported. Hue does not automatically connect to one of the Livy servers. You must specify the Livy server in the Hue Advanced Configuration Snippet as follows:

```
[desktop]
[spark]
livy_server_url=http(s)://[***LIVY-FOR-SPARK3-SERVER-HOST***]:
[***LIVY-FOR-SPARK3-SERVER-PORT***]
```

Moreover, you may see the following error in Hue when you submit a SparkSQL query: Expecting value: line 2 column 1 (char 1). This happens when the Livy server does not respond to the request from Hue.

Specify all different Livy servers in the `livy_server_url` property one at a time and use the one which does not cause the issue.

#### **CDPD-18491: PySpark and SparkSQL are not supported with Livy in Hue**

Hue does not support configuring and using PySpark and SparkSQL with Livy in CDP Private Cloud Base.

#### **Importing and exporting Oozie workflows across clusters and between different CDH versions is not supported**

You can export Oozie workflows, schedules, and bundles from Hue and import them only within the same cluster if the cluster is unchanged. You can migrate bundle and coordinator jobs with their workflows only if their arguments have not changed between the old and the new cluster. For example, hostnames, NameNode, Resource Manager names, YARN queue names, and all the other parameters defined in the workflow.xml and job.properties files.

Using the import-export feature to migrate data between clusters is not recommended. To migrate data between different versions of CDH, for example, from CDH 5 to CDP 7, you must take the dump of the Hue database on the old cluster, restore it on the new cluster, and set up the database in the new environment. Also, the authentication method on the old and the new cluster should be the same because the Oozie workflows are tied to a user ID, and the exact user ID needs to be present in the new environment so that when a user logs into Hue, they can access their respective workflows.



**Note:** Migrating Oozie workflows from HDP clusters is not supported.

### Known Issues in Kudu

Learn about the known issues in Kudu, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6

There are no new known issues identified in this release.

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6

**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.

Workaround: Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

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.

Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

### Unsupported features

Kudu HMS Sync is disabled and is not yet supported

### Known Issues in Spark

Learn about the known issues in Spark the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6.

### Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6

The following known issues were identified in this release:

### Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6

**CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files**

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the RuntimeWarning: Failed to add file message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

#### **CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress**

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

#### **CDPD-67517: Spark3 tests are failing if /tmp is mounted as NOEXEC.**

Map the tmpdir to a writable path in spark3-conf/spark-defaults.conf using the following steps:

1. Go to the Cloudera Manager.
2. Go to the Spark service.
3. Click the Configuration tab.
4. Select Scope Gateway .
5. Select Category Advanced .
6. Locate the Spark Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-defaults.conf\_client\_config\_safety\_valve property.
7. Map the tmpdir to a writable path:

```
spark.driver.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
spark.executor.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
```

8. Enter a Reason for change, and then click Save Changes to commit the changes.
9. Deploy the client configuration.

#### **CDPD-23817: In the upgraded Cluster, the permission of /tmp/spark is restricted due to the HDP configuration hive.exec.scratchdir=/tmp/spark.**

If you are using the /tmp/spark directory in the CDP cluster, you must provide the required additional Policies/ACL permissions.

#### **CDPD-22670 and CDPD-23103: There are two configurations in Spark, "Atlas dependency" and "spark\_lineage\_enabled", which are conflicted. The issue is when Atlas dependency is turned off but spark\_lineage\_enabled is turned on.**

Run Spark application, Spark will log some error message and cannot continue. That can be restored by correcting the configurations and restarting Spark component with distributing client configurations.

**CDPD-23007: Mismatch in the Spark Default DB Location.** In HDP 3.1.5, hive\_db entities have one attribute - 'location' which is configured to the '/managed' path. In fresh install of CDP 7.1.7, hive\_db entities now have 2 attributes 'location' configured to '/external' path and 'managedLocation' configured to '/managed' path. In the AM2CM migration (HDP 3.1.5 -> CDP 7.1.7), the 'location' attribute from hive\_db entities in HDP 3.1.5 comes unaltered to CDP 7.1.7 and hence maps to '/managed' path.

This issue arises only if you are upgrading from HDP 3.1.5 to CDP 7.1.7. If you are performing a fresh install of CDP 7.1.7, you can ignore this issue.

None

#### **CDPD-217: The Apache Spark connector is not supported**

The old *Apache Spark - Apache HBase Connector* (shc) is not supported in CDP releases.

Use the new HBase-Spark connector shipped in CDP release.

#### **CDPD-3038: Launching pyspark displays several HiveConf warning messages**

When pyspark starts, several Hive configuration warning messages are displayed, similar to the following:

```
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of name hive.meta
store.runworker.in does not exist
 23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of
name hive.masking.algo does not exist
 23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of
name hive.metastore.runworker.in does not exist
 23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of na
me hive.masking.algo does not exist
```

These errors can be safely ignored.

#### **Known Issues in Sqoop**

Learn about the known issues in Sqoop, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6**

There are no new known issues identified in this release.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6**

##### **CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version**

If you have upgraded to CDP PvC Base 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to CDP PvC Base 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see [Troubleshooting Apache Sqoop issues](#).

##### **CDPD-44431: Using direct mode causes problems**

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the --direct option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the `sqoop.enable.deprecated.direct` property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through `-Dsqoop.enable.deprecated.direct=true`.

#### **CDPD-3089: Avro, S3, and HCat do not work together properly**

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

#### **Parquet columns inadvertently renamed**

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.

#### **PARQUET-99: Importing Parquet files might cause out-of-memory (OOM) errors**

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.

#### **Known Issues in Cruise Control**

Learn about the known issues in Cruise Control, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6**

There are no new known issues identified in this release.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6**

##### **Rebalancing with Cruise Control does not work due to the metric reporter failing to report the CPU usage metric**

On the Kafka broker, the Cruise control metric reporter plugin may fail to report the CPU usage metric.

If the CPU usage metric is not reported, the `numValidWindows` in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.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 the [Moving a Host Between Clusters](#).



**Note:** Cluster nodes affected by this issue are not displayed as unhealthy.

#### **Known Issues in Kafka**

Learn about the known issues in Kafka, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6.

## Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6

There are no new known issues identified in this release.

## Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6

### CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

### OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL\_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker) in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Save Changes > Restart SMM.

### The `offsets.topic.replication.factor` property must be less than or equal to the number of live brokers

The `offsets.topic.replication.factor` broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a `GROUP_COORDINATOR_NOT_AVAILABLE` error until the cluster size meets this replication factor requirement.

None

### Requests fail when sending to a nonexistent topic with `auto.create.topics.enable` set to true

The first few produce requests fail when sending to a nonexistent topic with `auto.create.topics.enable` set to true.

Increase the number of retries in the producer configuration setting retries.

### Performance degradation when SSL is enabled

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with `ssl.secure.random.implementation = SHA1PRNG`. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

### OPSAPS-43236: Kafka garbage collection logs are written to the process directory

By default Kafka garbage collection logs are written to the agent process directory. Changing the default path for these log files is currently unsupported.

None

#### **RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure**

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transactional method because we are in an error state
 at org.apache.kafka.clients.producer.internals.TransactionManager.maybeFailWithError(TransactionManager.java:1125)
 at org.apache.kafka.clients.producer.internals.TransactionManager.maybeAddPartition(TransactionManager.java:442)
 at org.apache.kafka.clients.producer.KafkaProducer.doSend(KafkaProducer.java:1000)
 at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:914)
 at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:800)
 .
 .
 .
 Caused by: org.apache.kafka.common.errors.ClusterAuthorizationException: Cluster authorization failed.
```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting `enable.idempotence` to `false`.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

#### **CDPD-49304: AvroConverter does not support composite default values**

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

#### **DBZ-4990: The Debezium Db2 Source connector does not support schema evolution**

The Debezium Db2 Source connector does not support the evolution (updates) of schemas. In addition, schema change events are not emitted to the schema change topic if there is a change in the schema of a table that is in capture mode. For more information, see [DBZ-4990](#).

None.

#### **CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression**

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy compression. The HDFS Stateless Sink connector is only affected if the Compression Codec or Compression Codec for Parquet properties are set to SNAPPY.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: org.apache.hadoop.util.NativeCodeLoader.buildSupportsSnappy()
```

```
Failed to write to HDFS due to java.lang.RuntimeException: native snappy library not available: this version of libhadoop was built without snappy support.
```

Download and deploy missing libraries.



**Important:** Ensure that you complete steps 1-11 on all Kafka Connect hosts. Additionally, ensure that the advanced configuration snippet in step 12 is configured for all Kafka Connect role instances.

1. Create the /opt/nativelibs directory.

```
mkdir /opt/nativelibs
```

2. Change the owner to kafka.

```
chown kafka:kafka /opt/nativelibs
```

3. Locate the directory containing the Hadoop native libraries and copy its contents to the directory you created.

```
cp /opt/cloudera/parcels/CDH/lib/hadoop/lib/native/* /opt/nativelibs
```

4. Verify that 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**

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 Schema Registry**

Learn about the known issues in Schema Registry, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6**

There are no new known issues identified in this release.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6**

##### **CDPD-40380: Authorization checking issue when Kerberos is disabled**

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the `ranger.plugin.schema-registry.service.name` field.
3. Replace `GENERATED_RANGER_SERVICE_NAME` with the actual name of the service.
4. Restart the Schema Registry service.

##### **CDPD-49304: AvroConverter does not support composite default values**

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

##### **CDPD-59015: Schema Registry does not create new versions of schemas even if the schema is changed**

Schema Registry uses a schema fingerprinting mechanism to differentiate between schemas. However, fingerprinting does not take into consideration the schema attributes of the field type. As a result, if you have two schemas where the only difference is that one has type attributes defined and the other does not, they will be considered identical by Schema Registry. For example, the following schemas are considered identical:

```
#Schema V1
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [{"name": "local_timestamp_micros_long", "type": "long"}]}

#Schema V2
```

```
{
 "type": "record",
 "name": "schema_name",
 "namespace": "ns",
 "fields": [
 {
 "name": "local_timestamp_micros_long",
 "type": {
 "type": "long",
 "logicalType": "local-timestamp-micros"
 }
 }
]
}
```

Notice that the only difference is that in the second schema, the `local_timestamp_micros_long` field has a logical type specified. In cases like this, the new version of the schema is not created, the initial version is used. This is true even if the data that is being produced has a new schema version. The ID of the first schema version is used and is put in the serialized record. The new schema version is not created.

This issue is common when using change data capture (CDC) connectors like the Debezium connectors. This is because CDC connectors create schemas with the logical type decimal based on the column type in the database schema. For example:

```
{
 "type": "record",
 "name": "schema_name",
 "namespace": "ns",
 "fields": [
 {
 "name": "database_column",
 "type": {
 "type": "bytes",
 "logicalType": "decimal",
 "precision": 64,
 "scale": 0
 }
 }
]
}
```

If the database schema changes (for example, the column type), it is possible that only scale changes, which is a schema attribute.

```
{
 "type": "record",
 "name": "schema_name",
 "namespace": "ns",
 "fields": [
 {
 "name": "database_column",
 "type": {
 "type": "bytes",
 "logicalType": "decimal",
 "precision": 64,
 "scale": 1
 }
 }
]
}
```

In this case, even though scale changed to 1, the first version of the schema is used where scale is 0. As a result, the data is consumed with the wrong scale.

Avoid using logical types or other attributes. Alternatively, ensure that there are no changes in the logical types or other attributes between schema versions.

#### **OPSAPS-68708: Schema Registry might fail to start if a load balancer address is specified in Ranger**

Schema Registry does not start if the address specified in the Load Balancer Address Ranger property does not end with a trailing slash (/).

Set the value of the `RANGER_REST_URL` Schema Registry environment variable to an address that includes a trailing slash.

1. In Cloudera Manager, select the Schema Registry service.
2. Go to Configuration.
3. Find the Schema Registry Server Environment Advanced Configuration Snippet (Safety Valve) property and add the following:

```
Key: RANGER_REST_URL
Value: [***RANGER REST API URL***]
```

Replace `[***RANGER REST API URL***]` with an address that can be used by Schema Registry to access Ranger. Ensure that the address ends with a trailing slash. For example: `http://ranger-1.cloudera.com:6182/`

4. Restart the Schema Registry service.

#### **OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade**

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

1. Access the Ranger Console (Ranger Admin web UI).
2. Click the `cm_atlas` resource-based service.
3. Add the `schemaregistry` user to the `all - *` policies.
4. Click `Manage Service Edit Service`.

5. Add the schemaregistry user to the default.policy.users property.

**OPSAPS-69317: Kafka Connect Rolling Restart Check fails if SSL Client authentication is required**

The rolling restart action does not work in Kafka Connect when the `ssl.client.auth` option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set `ssl.client.auth` to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

**Known Issues in Streams Messaging Manager**

Learn about the known issues in Streams Messaging Manager, the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6.

**Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6**

There are no new known issues identified in this release.

**Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6****OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager**

Cloudera Manager does not display a Log Files menu for SMM UI role (and SMM UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by SMM UI is not supported by Cloudera Manager.

View the SMM UI logs on the host.

**OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners**

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL\_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in SMM safety valve in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Add the following value for bootstrap servers>Save Changes > Restart SMM.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

**CDPD-70313: KNOX does not send Authentication header on FIPS configuration**

KNOX does not send neither the Authentication header nor the `hadoop.auth` cookie. Because of this, the SMM UI returns an HTTP 401 response and sets the `www-authenticate: Negotiate` header. After this, KNOX still does not send the Authentication header. This results in the SMM UI being inaccessible through Knox.

You can access the SMM UI directly in Cloudera Manager at Clusters SMM Streams Messaging Manager Web UI and log in using the Kerberos username and password.

**CDPD-82560: Streams Messaging Manager known unprotected endpoint**

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**

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**

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.

#### **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 in Cloudera Runtime 7.1.9 SP1 CHF 6.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6**

There are no new known issues identified in this release.

#### **Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6**

##### **CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic**

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

##### **CDPD-11079: Blacklisted topics appear in the list of replicated topics**

If a topic was originally replicated but was later excluded for replication, 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, this topic will be included in the response. Additionally, the excluded topic will also be visible in the SMM UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

##### **CDPD-30275: SRM may automatically re-create deleted topics on target clusters**

If `auto.create.topics.enable` is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with `srm-control`. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1]
```

2. Wait until SRM is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

#### **Known Issues in Zeppelin**

Learn about the known issues in Zeppelin the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6.

#### **Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6**

The following known issues were identified in this release:

**Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6****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.

**BUG-125263: Zeppelin service move fails on clusters upgraded from HDP3.1.5**

Resolve the circular symlink issue on the Zeppelin node by linking the conf directory to a new directory under /etc/zeppelin:

- # mkdir -p /etc/zeppelin/<version>/0
- # rm /usr/hdp/<version>/zeppelin/conf
- # ln -s /etc/zeppelin/<version>/0 /usr/hdp/<version>/zeppelin/conf

Where version is the HDP version. For example, 7.1.x-yyy. Restart the Zeppelin server in Ambari.

**CDPD-3090: Due to a configuration typo, functionality involving notebook repositories does not work**

Due to a missing closing brace, access to the notebook repositories API is blocked by default.

From the CDP Management Console, go to Cloudera Manager for the cluster running Zeppelin. On the Zeppelin configuration page (Zeppelin serviceConfiguration), enter shiro urls in the Search field, and then add the missing closing brace to the notebook-repositories URL, as follows:

```
/api/notebook-repositories/** = authc, roles[{{zeppelin_admin_group}}]
```

Click Save Changes, and restart the Zeppelin service.

**CDPD-2406: Logout button does not work**

Clicking the Logout button in the Zeppelin UI logs you out, but then immediately logs you back in using SSO.

Close the browser.

**OPSAPS-59802: Zeppelin and Livy roles should be co-located on the same host.**

When installing or upgrading to CDP Private Cloud Base, you must co-locate all Zeppelin and Livy roles on the same cluster host due to an issue with certificate generation.

**Known Issues in ZooKeeper**

Learn about the known issues in ZooKeeper the impact or changes to the functionality, and the workaround in Cloudera Runtime 7.1.9 SP1 CHF 6.

**Known issues identified in Cloudera Runtime 7.1.9 SP1 CHF 6**

There are no new known issues identified in this release.

**Known issues identified before Cloudera Runtime 7.1.9 SP1 CHF 6**

There are no new known issues identified in the previous releases.

**Cumulative hotfix 5**

Know more about the cumulative hotfix 5 for CDP 7.1.9 SP1. This cumulative hotfix was released on 20 Feb, 2025.

**Cloudera Runtime 7.1.9.1032 (Cumulative Hotfix 5) download URL:****Parcel Repository Location**

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1032/parcels/
```

## What's new in Cloudera Runtime 7.1.9 SP1 CHF 5

Understand the functionalities and improvements to features of components in Cloudera Runtime 7.1.9 SP1 CHF 5.

### What's new in Hue

#### File extension restrictions for Hue file uploads

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](#).

### Fixed issues in 7.1.9 SP1 CHF 5

Know more about the cumulative hotfix 5 for 7.1.9 SP1. This cumulative hotfix was released on 20 Feb, 2025.



**Note:** Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p1032.62597146.

#### CDPD-76634: RMS server throws ConcurrentModificationException

The original ConcurrentModificationException is likely thrown when the resource-mappings are modified in response to changes in the Hive metadata while they are 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-75981: Default native ACL limits to user and user's primary group

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-77877: Exclude guava from arrow-vector

Fixed an issue where a transitive dependency would force custom applications to be built with a particular version of guava.

#### CDPD-59896: Hue fails to clean up documents if Oozie is on the blocked list of services

Earlier, when you ran the desktop\_document\_cleanup script to clean old documents, and if you added Oozie to the blocked list of apps in the Hue Advanced Configuration Snippet, the script failed. To run the script successfully, you had to remove the Oozie service from the app\_blacklist property manually. This issue has been resolved. You no longer need to manually remove Oozie from the blocked list of apps.

#### CDPD-62397: Improved default editor selection in Hue configuration

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 has now been resolved, allowing you to specify a default editor in the Hue configuration file for all users.

#### CDPD-75286: Spark History UI - StreamConstraintsException: String length exceeds the maximum length

Fixed an issue with Jackson to allow unlimited json string length in Spark event logs.

#### Zeppelin Livy interpreter paragraph outputs are being misplaced

Fixed an issue where Zeppelin paragraph outputs were being misplaced. Users are now able to use pyspark to run os.system commands without issues.

**CDPD-78260: Partition value mismatch during INSERT event processing**

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-77922: Prometheus metric name incompatibility with JDK11 and JDK17**

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-73427: Upgrade Spring Framework to resolve vulnerabilities**

CVE-2024-38808: Vulnerable to resource exhaustion through user-supplied SpEL expressions.

CVE-2024-38809: DoS attack possible due to improper ETag validation in HTTP requests.

Upgraded Spring Framework to versions 6.1.12, 6.0.23, or 5.3.39 to mitigate these issues.

**CDPD-75665: Importing a table generates a DDL with an incorrect location**

When creating a table using the IMPORT command, an external imported table can have this partitions located under the managed warehouse directory, which violates the metastore.warehouse.external.dir and metastore.warehouse.dir that intend to host different type of tables.

The fix ensures that the table path is fetched from HMS before importing partitions, ensuring that all data is copied to the correct location in the newly imported table.

Apache Jira: [HIVE-28580](#)

**CDPD-74658: Invalid plan generated for IN subquery with multiple correlations**

Hive was generating incorrect query plans for IN subqueries with multiple correlations, causing errors due to missing column references.

This issue is now fixed.

Apache Jira: [HIVE-24999](#)

**CDPD-76983 CDPD-76982: Removed redundant code**

This issue is now fixed.

Apache Jira: [HIVE-24412](#)

Apache Jira: [HIVE-18725](#)

**CDPD-76985: Optimize scalar query rewrite to avoid unnecessary joins**

Hive planner adds an unnecessary join for scalar subqueries that return at most one row.

The subquery removal rule now checks row count estimates and skips adding the extra join.

Apache Jira: [HIVE-21537](#)

**CDPD-76980: Support EXISTS/NOT EXISTS correlated subqueries with aggregates**

Hive failed to handle correlated EXISTS/NOT EXISTS subqueries containing aggregate functions without a GROUP BY clause, leading to query compilation errors.

This functionality is now supported.

Apache Jira: [HIVE-15757](#)

**CDPD-76978: Improve error handling for invalid column references in subqueries**

Hive returned a misleading error when a subquery referenced a non-existing column, making it unclear that the issue was a missing column

Removed HiveSubQRemoveRelBuilder and used RelBuilder.

Apache Jira: [HIVE-24685](#)

**CDPD-74657: Incorrect results when using COALESCE in correlated subqueries**

Queries with a correlated subquery using the COALESCE function in the correlation predicate returned incorrect results.

The query execution logic was updated to correctly apply COALESCE in correlated subqueries, ensuring that all matching rows are included in the result set.

Apache Jira: [HIVE-24957](#)

**CDPD-76984: Incorrect query results due to merging TableScan operators with different Dynamic Partition Pruning parents**

The shared work optimizer merged TableScan operators that had different Dynamic Partition Pruning (DPP) parents. This caused missing filters in the query plan, leading to incorrect results.

The issue is now fixed.

Apache Jira: [HIVE-26968](#)

**CDPD-78077: Kafka Connect connectors might fail during deployment or while being moved to another node**

Kafka Connect connectors no longer fail due to incompletely unpacked dependencies while being deployed or moved to another node.

**CDPD-76135: Schema check database connection throws PSQLErrorException**

The schema check now completes successfully during Oozie startup for secure database connections or when custom connection properties are used, preventing warnings from appearing in the logs.

**CDPD-76604: CDEP fails to install CDP Private Cloud Base as Ozone does not start up. The OZONE-1 has failed in startup so CDEP hasn't executed next set of commands. error is displayed.**

This issue is resolved by fixing the race condition in creation and deletion of rocksdbjni native library.

**COMPX-18909: NodeManager can be marked unhealthy if an application is killed**

A logic was introduced to mark the NodeManager unhealthy if the container-executor encountered a unrecoverable/configuration error, that is, when the container-executor script was missing. If the application was killed just before when one of the containers tried to access the localizer syslog file an IOException occurred due to the deleted app log directory, and then the NodeManager was marked unhealthy incorrectly.

This issue is now resolved and the error checking is more specific and the false positive NodeManager markings are now removed.

Apache Jira: [YARN-11753](#)

**CDPD-49702: NodeManager must be shut down when the program /var/lib/yarn-ce/bin/container-executor cannot be run**

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](#)

**CDPD-74510: Migration status is not updated on file while migration is in progress**



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.

**COMPX-18663: Cgroup v2 support must fall back to v1 when there are no v2 controllers**

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](#)

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

- [CVE-2021-41973](#) - Apache Mina
- [CVE-2024-22243](#) - Spring Framework
- [CVE-2024-22259](#) - Spring Framework
- [CVE-2024-22262](#) - Spring Framework
- [CVE-2024-38808](#) - Spring Framework
- [CVE-2024-38809](#) - Spring Framework
- [CVE-2024-38816](#) - Spring Framework
- [CVE-2024-38819](#) - Spring Framework
- [CVE-2024-38820](#) - Spring Framework
- [CVE-2024-38821](#) - Spring Framework



**Note:** Starting with CDP Private Cloud Base 7.1.9 SP1 CHF5, Cloudera introduced an internal fork of Java Spring Framework. For more details, see [Cloudera's new internal spring framework](#).

**Known issues in 7.1.9 SP1 CHF 5**

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.1.9 SP1 CHF 5.

The list of known issues for the Runtime release 7.1.9 SP1 CHF 5 includes the known issues from Runtime release 7.1.9 SP1. For more information, see [Known issues in Cloudera Runtime 7.1.9 SP1](#).

**CDPD-49323: INSERT statement does not respect Ranger policies for HDFS**

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to `cm_hdfs` -> all-path for the user `impala`, inserting into a table whose HDFS POSIX permissions happen to exclude `impala` access will result in `AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXX XXXXXX`

**Apache Jira:** [IMPALA-11871](#)

**CDPD-78069: Oozie action configuration's Java options are not applied due to CDPD-60551**

Java options configured in workflow xml's action configurations such as `yarn.app.mapreduce.am.command-opts` and `mapreduce.map.java.opts` are not applied to the Oozie Launcher AM's JVM.

1. Go to Cloudera Manager Oozie Configuration .
2. Search for Oozie Server Advanced Configuration Snippet (Safety Valve) for `oozie-site.xml`.

3. In the Oozie Server Advanced Configuration Snippet (Safety Valve) for `oozie-site.xml`, add the following line:

```
<NAME>oozie.LauncherConfigurationInjector.hadoop.search.properties</name><VALUE> </VALUE>
```



**Important:** The value is a single `/space/` character.

4. Restart Oozie.

#### **CDPD-77054: Exporting shell entities can cause the import to fail**

For Apache Hive tables created with Apache Spark, there may be shell entities created in Apache Atlas. If these entities are in the export zip file before they are resolved, the import will fail.

Workaround: None

#### **CDPD-75742: Upgrade commons-lang3 to 3.13.0**

The Commons-text 1.11.0 in Spark calls `Range.of(..)` method of `commons-lang3` which is not available in the version 3.12.0, resulting in Oozie-Spark job failure.

Upgrade to `commons-lang3` version 3.13.0 which contains the required method.

#### **CDPD-77399: HBase fails to register the servlet metrics and throws `ClassNotFoundException`: `org.apache.hadoop.metrics.MetricsServlet`**

The `MetricsServlet` class is a Hadoop 2-based metric servlet unavailable in Hadoop 3 deployments.

Workaround: Ignore this WARN log message during HBase Master and RegionServer startup.

Apache Issue: [HBASE-28315](#)

#### **CDPD-70450: Impala SQL queries that include the “WITH” clause should populate lineage in Atlas**

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-76035: Resource lookup for Atlas service is failing**

Once the Atlas configuration snippet `atlas.authentication.method.file` is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

#### **CDPD-77738: Atlas hook authorization issue causing `HiveCreateSysDb` timeout**

Atlas hook authorization error causes `HiveCreateSysDb` command to time out due to repeated retries.

None

#### **CDPD-79160: NPE while deleting `BusinessMetadata`**

If business metadata is created without adding any applicable types, a `NullPointerException` is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

#### **CDPD-61527: On FIPS cluster, Knox-Impala connection is failing with SSL error code 5**

The Knox-Impala connection is failing with SSL error code 5 on FIPS clusters. This prevents accessing Impala through Knox.

#### **CDPD-78656: Health test for Knox fails if the `gateway.client.auth.needed = true` is set**

The health test for Knox Gateway fails if the `gateway.client.auth.needed` parameter is set to true. Environments using the "curl" call are impacted. The curl call from CM is not specifying any certificate (store) while Knox is configured to require one.

To avoid the startup failure, use `gateway.client.auth.wanted = true`.

#### **CDPD-77911: Missing Log4j Redactor dependency**

The missing `org.cloudera.logredactor` dependency can result impossible Log4j errors during startup. Additionally, it can lead to unwanted data leak because of incorrect log redaction.

#### **CDPD-80921: Without permission for one glossary, /glossary call throws exception**

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-82056: UI: when server response date fields as '0', UI shows as current time**

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](#)

#### **CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export**

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**

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](#)

#### **CDPD-66938: [Analyze] [Atlas] test\_time\_range tests fail**

When the Apache Atlas server is running on a node which has time zone other than UTC, there might be a time of day when the search results might differ if the relative CreateTime date range filters of TODAY, YESTERDAY, etc. are used. 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.

Use explicit date range filters instead of using relative date range filters, such as, TODAY, YESTERDAY.

#### **CDPD-76789: Creating tag with name description throws java.lang.ClassCastException**

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails**

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

#### **CDPD-70321: Atlas Parallel import is failing with various errors**

During a parallel import-export activity with six iceberg table policies with exportOption as db1.\* for all six exports, all import fail after the exports.

#### **CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen**

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

#### **Crash in Navigator Encrypt kernel module if --profile-file= option is used with ACLs**

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

#### **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.

#### **UnsupportedClassVersionError**

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.

#### **CDPD-87670: [Atlas UI] Apache Atlas Glossary becomes unresponsive when the page size is set to 50**

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn**

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a `LockingException` (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the `PermanentLockingException`.

#### **CDPD-79099: Unable to import hive tables using import script in Public Cloud**

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-88230: Bulk import operations currently do not construct lineage as expected**

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**

Bulk import operations do not create audits for the imported entities.

**CDPD-45370: Cartesian join fails with an empty table**

When the configuration `hive.tez.cartesian-product.enabled` is set to true, query fails if one of the tables involved in a Cartesian join has zero rows. This results in an `IllegalArgumentException` with the message `Number of items is 0. Should be positive`.

To address this issue, set `hive.tez.cartesian-product.enabled` to false before running the query.

Apache Jira: [\[HIVE-26582\]](#)

**Technical Service Bulletins****TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager**

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.

**Knowledge article**

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

**CDPD-84502: Advanced Search is not working properly in Atlas UI**

The Advance Search in Atlas does not return any results when selecting the Typedef option. The API is not triggered for each relation.

None

**Cumulative hotfix 4**

Know more about the cumulative hotfix 4 for CDP 7.1.9 SP1. This cumulative hotfix was released on Jan 13, 2025.

**Cloudera Runtime 7.1.9.1023 (Cumulative Hotfix 4) download URL:****Parcel Repository Location**

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1023/parcels/
```

**Fixed issues in 7.1.9 SP1 CHF 4**

Know more about the cumulative hotfix 4 for 7.1.9 SP1. This cumulative hotfix was released on 10 Jan, 2025.



**Note:** Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p1023.60818430.

**COMPX-18033: Fixed the order of updating CPU controls with cgroup v1**

This fix updates `cpu.cfs_period_us` 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](#)

**COMPX-17463: Setting maximum-application-lifetime using AQCv2 templates did apply to the first submitted application**

Previously, when the maximum-application-lifetime property was set using the AQC v2 templates (`yarn.scheduler.capacity.root.test.auto-queue-creation-v2.template.maximum-application-lifetime`), the property did not apply to the first submitted application, and was only applied to the subsequent ones. This issue is now resolved.

Apache Jira: [YARN-11708](#)

**CDPD-76503: Query with UNION operators fails due to OutOfMemoryError**

Queries containing multiple UNION operators failed with a Java heap space `OutOfMemoryError` when executing or generating the explain plan.

The issue was addressed by optimizing how metadata is managed for UNION operators, preventing excessive memory usage during query execution, or explaining plan generation. This ensures queries with UNION operators run without causing memory errors.

Apache Jira: [HIVE-25220](#)

**CDPD-76378: fs.s3a.connection.expect.continue controls 100 CONTINUE behavior**

Added `fs.s3a.connection.expect.continue` option to control whether or not a PUT request to the S3 store sets the "Expect: 100-continue" header, and waits for the 100 CONTINUE response before uploading any data. The configuration parameter is set to true by default.

Apache Jira: [HADOOP-19317](#)

**CDPD-71872: Failed table loads not retried after metastore recovery**

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-75891: Missing partitions in COPY TESTCASE with LocalCatalog mode**

In LocalCatalog mode, COPY TESTCASE did not include partition and file metadata, only exporting table schema and column stats. Importing the test case caused partition ID conflicts, leading to incomplete or incorrect results.

- Partition and file metadata are now correctly included during export.
- Unique partition IDs are generated during import to avoid conflicts.
- Logic for handling partitions and unpartitioned tables during import has been fixed.

These updates ensure that COPY TESTCASE works reliably in LocalCatalog mode.

Apache Jira: [IMPALA-13340](#)

**CDPD-76186: Ozone does not support snapshots on linked buckets. This can impact Replication Manager jobs running on linked W3 buckets.**

This issue is resolved by supporting Ozone snapshots on linked buckets. However, if you have created any snapshots on the linked buckets from the versions lower than 7.1.9 SP1 CHF4, you must delete the snapshots on the linked buckets by running the `ozone sh snapshot delete`

<vol>/<linked bucket name> <snapshot name> command before upgrading to 7.1.9 SP1 CHF4.

Apache Jira: [HDDS-11705](#)

#### **CDPD-75721: LOAD DATA INPATH fails with ICEBERG table in Impala**

If you are using Impala to run a LOAD DATA INPATH statement on an Iceberg table, the statement fails if the target table is fully qualified.

This issue is now resolved.

#### **CDPD-76136: Schema resolution does not work for migrated partitioned Iceberg tables that have complex types**

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-75422/CDPD-76137: Impala should ignore character case of Iceberg schema elements**

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-76091: Duplicate entries in TXN\_WRITE\_NOTIFICATION\_LOG causing replication failures**

Oracle's treatment of empty strings as NULL led to duplicate entries in `TXN_WRITE_NOTIFICATION_LOG` during Hive ACID transactions with MERGE statements. This caused mismatches in `_files` and `_dumpmetadata` files, resulting in incremental LOAD failures and blocking replication.

The query logic was updated to exclude partition predicates for non-partitioned tables, ensuring consistent dump files and resolving replication issues.

#### **CDPD-75941: Spark Iceberg vectorized Parquet read of decimal column is incorrect**

Fixed an issue with Spark reading Iceberg Parquet data files with a decimal column incorrectly when vectorized read was enabled.

#### **CDPD-75925: Append realm of the Kerberos ticket to the navencrypt-gen-keytab command**

It is now necessary to add the realm of the kerberos ticket to the "navencrypt-gen-keytab" command when creating the ticket. For example, `navencrypt-gen-keytab https://abc-1.cloudera.com:7183 admin ABC.CLOUDERA.COM`.

#### **CDPD-75500: Correct Handling of ResourceTrie Node Removal to Ensure Accurate Policy Evaluation with Policy-Deltas Enabled**

When policy-deltas are enabled, if two policies share a common subset of resources and are assigned to the same user (either directly or through groups), modifying any attribute of one policy (such as its name, resources, or users) causes only the modified policy to be considered during access evaluation. This issue persisted until the underlying service (for example, HDFS) was restarted.

The root cause was that the ResourceTrie node referencing the modified policy evaluator was removed, even if it still contained wildcard evaluators. The fix ensures that the self node is removed



from the ResourceTrie only if it has no children, no evaluators, and no wildcard evaluators. As a result, both policies are now correctly considered during access evaluation after any modification.

**CDPD-75488: Fails to update to Phoenix 5.1.3 with PhoenixNonRetryableRuntimeException**

While updating to the Phoenix 5.1.3 version, the PhoenixNonRetryableRuntimeException is thrown which states that the `org.apache.phoenix.query.DefaultGuidePostsCacheFactory` class cannot be loaded or instantiated from SquirrelSQL

Apache Jira: [PHOENIX-7290](#)

**CDPD-75383: Ranger Audit Filter for the HBase service is not working as expected**

From the UI perspective, while creating services configuring the Audit Filter resources, incorrect values were set for the 'Include' and 'Exclude' flag setup. This issue has now been resolved and is tracked in this Jira ticket

Apache Jira: [RANGER-4983](#)

**CDPD-75008: Performance fixes for Ozone plugin**

Fixes the performance issues observed while evaluating policies for multi-level resources.  
RANGER-4893 : Improves policy evaluation for multi level resource hierarchies. RANGER-4922 : Reduces time to find tags associated with multi-level resources.

**CDPD-74821: Incorrect query results when running concurrent Hive queries**

Running multiple concurrent Hive queries through Beeline could produce incorrect results. This issue occurred when the `hive.query.result.cached.enabled` setting was set to false. The queries fetched results from other parallel queries instead of their own, leading to inconsistencies in workflows.

The issue has been resolved by ensuring that files are set in a thread-safe manner within `HiveSequenceFileInputFormat`.

Apache Jira: [HIVE-28530](#)

**CDPD-74730: Backport SPARK-46239: Hide the Jetty server's version**

[SPARK-46239] Hide the Jetty server's version.

Apache Jira: [SPARK-46239](#)

**CDPD-74483: Ozone replication from source to destination fails when `ozone.client.bytes.per.checksum` is set to 16 KB.**

When `ozone.client.bytes.per.checksum` is set to a lower value, say 16 KB, the parity checksum calculation during the validation phase is wrong and displays the `IllegalArgumentException` error on the client side.

This issue is now resolved by setting `ozone.client.bytes.per.checksum` to 1 MB.

**CDPD-72174: During the `FileSizeCountTask`, the `ClassCastException` error is displayed in the Recon server where `RepeatedOmKeyInfo` is incorrectly cast to `OmKeyInfo` causing task processing to fail.**

This issue is resolved by implementing a safeguard that checks for value mismatches and ignores such events.

Apache Jira: [HDDS-11187](#)

**CDPD-71931: Ranger - Upgrade commons-compress to 1.26.0 due to CVE-2024-25710 and CVE-2024-26308**

Upgraded commons-compress to 1.26.0 due to CVE-2024-25710 and CVE-2024-26308

Apache Jira: [RANGER-4729](#)

**CDPD-71411: Atlas Hbase import fails with `jackson-databind` exception**



Updated jackson-annotations to match jackson-databind (2.15.0) to resolve NoSuchFieldError during HBase import caused by version mismatch (older 2.12.7 on classpath).

**CDPD-68425: Oozie should validate the backend database identifiers in case-insensitive way**

Resolved the issue where Oozie incorrectly reported missing tables due to issues with table name casing.

**CDPD-62726: Atlas 'updateTime' parameter is not updated when business metadata, labels is edited.**

The change, fixed the issue when the updateTime system attribute was not updated for business metadata and labels when added and removed from an entity. This caused the customer significant issues with audit and accountability of user actions in this governance tool, and severely limits its use within the organisation, as the business metadata functionality is used for all data.

Also, this helped the incremental import export where updated actions for business metadata and labels were performed on the entity.

Apache Jira: [ATLAS-4802](#)

**CDPD-44060: Stale locks preventing materialized view rebuilds**

The ALTER MATERIALIZED VIEW <view\_name> REBUILD command failed with the following error:

```
Error: Error while compiling statement: FAILED: SemanticException
org.apache.hadoop.hive.ql.parse.SemanticException: Another process
is rebuilding the materialized view <view_name>
```

This occurred because stale lock entries in the backend materialization\_rebuild\_locks table were not removed when the rebuild operation failed or was interrupted. These stale locks prevented subsequent rebuild attempts for the same materialized view.

The issue was addressed by ensuring that stale lock entries are removed from the materialization\_rebuild\_locks table even when a rebuild transaction is aborted. This ensures that subsequent rebuild commands for the same materialized view can proceed without errors.

Apache Jira: [HIVE-28416](#)

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

- [CVE-2021-47621](#) - Classgraph

**Known issues in 7.1.9 SP1 CHF 4**

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.1.9 SP1 CHF 4.

The list of known issues for the Runtime release 7.1.9 SP1 CHF 4 includes the known issues from Runtime release 7.1.9 SP1. For more information, see [Known issues in Cloudera Runtime 7.1.9 SP1](#).

**CDPD-70450: Impala SQL queries that include the “WITH” clause should populate lineage in Atlas**

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-75665: Importing a table generates a DDL with an incorrect location**

When creating a table using the IMPORT command, this table's partitions could point to an incorrect location, that is, an external imported table can have his partitions located under the managed warehouse directory, which violates the metastore.warehouse.external.dir and metastore.warehouse.dir that intend to host different type of tables.

Apache Jira: [HIVE-28580](#)

None

**CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout**

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

**COMPX-18909:YARN NodeManagers potentially getting into unhealthy state caused by YARN-11709**

A logic was introduced to mark the NodeManager unhealthy if the container-executor encounters an unrecoverable or configuration error. For example, if the container-executor script is missing from the configured place. There is an edge-case in this behaviour that, if the application is killed just before one of the containers is trying to access the localizer syslog file, an IOException occurs because the NodeManager has already deleted the application log directory. This IOException is then caught by the logic of CDPD-49702/YARN-11709, and the NodeManager is marked unhealthy incorrectly.

Restart the unhealthy NodeManager.

**OPSAPS-72509, CDPD-32440: Hive metadata transfer to GCS fails with ClassNotFoundException**

Hive replication policies from an on-premises cluster to cloud fails during the “Transfer Metadata Files” step if the following conditions are true:

- the target is a GCS Data Lake
- the source Cloudera Manager version is 7.11.3 CHF7, 7.11.3 CHF8, 7.11.3 CHF9, 7.11.3 CHF9.1, 7.11.3 CHF10, or 7.11.3 CHF11

This is because the fs.gs.delegation.token.binding property is already defined in the configuration and cannot be unset to disable the delegation tokens in the cloud connector service.

None

**CDPD-75742: Upgrade commons-lang3 to 3.13.0**

The Commons-text 1.11.0 in Spark calls Range.of(..) method of commons-lang3 which is not available in the version 3.12.0, resulting in Oozie-Spark job failure.

Upgrade to commons-lang3 version 3.13.0 which contains the required method.

**CDPD-77399: HBase fails to register the servlet metrics and throws ClassNotFoundException: org.apache.hadoop.metrics.MetricsServlet**

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](#)

**CDPD-76035: Resource lookup for Atlas service is failing**

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

**CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout**

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

**CDPD-79160: NPE while deleting BusinessMetadata**

If business metadata is created without adding any applicable types, a NullPointerException is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

**CDPD-61527: On FIPS cluster, Knox-Impala connection is failing with SSL error code 5**

The Knox-Impala connection is failing with SSL error code 5 on FIPS clusters. This prevents accessing Impala through Knox.

**CDPD-78656: Health test for Knox fails if the `gateway.client.auth.needed = true` is set**

The health test for Knox Gateway fails if the `gateway.client.auth.needed` parameter is set to true. Environments using the "curl" call are impacted. The curl call from CM is not specifying any certificate (store) while Knox is configured to require one.

To avoid the startup failure, use `gateway.client.auth.wanted = true`.

**CDPD-77911: Missing Log4j Redactor dependency**

The missing `org.cloudera.logredactor` dependency can result impossible Log4j errors during startup. Additionally, it can lead to unwanted data leak because of incorrect log redaction.

**CDPD-80921: Without permission for one glossary, /glossary call throws exception**

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-82056: UI: when server response date fields as '0', UI shows as current time**

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](#)

**CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export**

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**

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](#)

**CDPD-66938: [Analyze] [Atlas] test\_time\_range tests fail**

When the Apache Atlas server is running on a node which has time zone other than UTC, there might be a time of day when the search results might differ if the relative CreateTime date range filters of TODAY, YESTERDAY, etc. are used. 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.

Use explicit date range filters instead of using relative date range filters, such as, TODAY, YESTERDAY.

**CDPD-84502: Advanced Search is not working properly in Atlas UI**

The Advance Search in Atlas does not return any results when selecting the Typedef option. The API is not triggered for each relation.

None

**CDPD-76789: Creating tag with name description throws `java.lang.ClassCastException`**

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails**

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

#### **CDPD-70321: Atlas Parallel import is failing with various errors**

During a parallel import-export activity with six iceberg table policies with exportOption as db1.\* for all six exports, all import fail after the exports.

#### **CDPD-50493: Data preview on database views in Hue**

The data preview functionality for views in Hue is disabled due to resource strains caused by selecting data from complex views built on multiple tables. This change addresses performance issues; Disabling the data preview functionality significantly impacted the workflows that rely on data preview for efficient validation and analysis in Hue.

When you try to view sample data from the Table Browser in Hue, Hue automatically runs the `SELECT * FROM <database>.<table> LIMIT 100` query from tables and views. To generate data from views, Hue fetched all rows and then displayed 100, based on the limit set in the query.

#### **CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeUpException: null" is seen**

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

#### **Crash in Navigator Encrypt kernel module if --profile-file= option is used with ACLs**

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

#### **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.

#### **UnsupportedClassVersionError**

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.

#### **CDPD-87670: [Atlas UI] Apache Atlas Glossary becomes unresponsive when the page size is set to 50**

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn**

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

**CDPD-79099: Unable to import hive tables using import script in Public Cloud**

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-88230: Bulk import operations currently do not construct lineage as expected**

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**

Bulk import operations do not create audits for the imported entities.

**CDPD-45370: Cartesian join fails with an empty table**

When the configuration hive.tez.cartesian-product.enabled is set to true, query fails if one of the tables involved in a Cartesian join has zero rows. This results in an IllegalArgumentException with the message Number of items is 0. Should be positive.

To address this issue, set hive.tez.cartesian-product.enabled to false before running the query.

Apache Jira: [\[HIVE-26582\]](#)

## Technical Service Bulletins

**TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager**

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.

## Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

## Cumulative hotfix 3

Know more about the cumulative hotfix 3 for 7.1.9 SP1. This cumulative hotfix was released on 11 Nov, 2024.

### Cloudera Runtime 7.1.9.1015 (Cumulative Hotfix 3) download URL:

#### Parcel Repository Location

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1015/parcels/
```

### What's new in Cloudera Runtime 7.1.9 SP1 CHF 3

Understand the functionalities and improvements to features of components in Cloudera Runtime 7.1.9 SP1 CHF 3.

#### Ranger

##### Support for SASL Bind in Ranger Usersync

Earlier, Usersync of Ranger supported Simple Bind to authenticate to a directory service. Now, Usersync of Ranger supports GSSAPI SASL Bind using Kerberos mechanism to authenticate to a directory service from Usersync. For more information on the configurations, see [Configure SASL Bind in Ranger Usersync](#).

#### Atlas

##### Ozone File System Support for Ranger Tagsync

In the Cloudera Private Cloud 7.1.9 SP1 CHF3 version, Apache Ranger Tagsync now supports Ozone file system. In this version, it is possible to apply tag based policies for OFS entities tagged in Atlas. Also, tag based access control is now applied recursively for subdirectories of a tagged O3FS key in Atlas. For more information, see [About Apache Ozone integration with Apache Atlas](#).

### Behavioral changes in Cloudera Runtime 7.1.9 SP1 CHF 3

You can review the changes in certain features or functionalities of components that have resulted in a change in behavior from the previously released version to this version of Cloudera Runtime 7.1.9 SP1 CHF 3.

#### Atlas

##### Summary:

The Exclude SubTypes and Exclude Sub-classifications filters were removed from the **Table** tab of entity details.

Previous behavior:

Previously, the Exclude SubTypes and Exclude Sub-classifications filters were available from the **Table** tab in entity details. There were no properties being passed to these filters when you visited the entity details of the page.

New behavior:

The two unused filter checkboxes Exclude SubTypes and Exclude Sub-classifications from the **Table** tab of entity detail page were removed.

##### Summary:

Special character validation was added to glossary, term and category names in Apache Atlas.

Previous behavior:

The special characters '@', '.', '<', '>' could be used in glossary, term and category name fields.

New behavior:

The special characters '@', '.', '<', '>' are no longer accepted in glossary, term and category name fields by the validation introduced. Avoid using these characters when creating glossary names, glossary terms and category names.

#### Summary:

Recursive tag based access control for Ozone keys, Apache Ranger Tagsync OFS support is introduced.

Previous behavior:

Ozone keys were not recursively checked for tag based policies. Tag based policy was only applicable for the directory level the tag is applied to. Additionally, Ranger Tagsync did not support OFS paths, thus tag-based policies did not work for OFS paths previously.

New behavior:

Ozone keys will now be recursively checked for tags and tag based policies. This means that tags applied for parent directory will be applicable to subdirectories too. While this is an expected behavior, this may cause behavior change in access control for customers already using tag based policies for Ozone keys.

The behavior change will be for new keys tagged in Atlas after upgrading.

Tag based policies will now work for OFS entities too.

### Ranger

#### 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 7.1.9 SP1 CHF3, 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 7.1.9 SP1 CHF3 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.

### Fixed issues in 7.1.9 SP1 CHF 3

Know more about the cumulative hotfix 3 for 7.1.9 SP1. This cumulative hotfix was released on 11 Nov, 2024.



**Note:** Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p1015.59425156.

#### CDPD-75272: Backport YARN-11709 to 7.1.9.1000

NodeManager is now shut down when it cannot run the `/var/lib/yarn-ce/bin/container-executor` program.

Apache Jira: [YARN-11709](#)

#### CDPD-75008: Performance fixes for Ozone plugin

Fixed the performance issues observed while evaluating policies for multi-level resources:

- Improves policy evaluation for multi level resource hierarchies (RANGER-4893).
- Reduces time to find tags associated with multi-level resources (RANGER-4922).

**CDPD-75001: Export/Import APIs : Unable to import shell entities**

Validation of mandatory attributes for shell entity has been removed

**CDPD-74686: [AUTOSYNC] DirectoryDeletion task ignored through ratis**

Previously, directory deletion task was ignored through Ratis, and there was deletion retry instead of actually deleting it. This issue is now resolved.

**CDPD-74685: [AUTOSYNC] Directory deletion get stuck having millions of directory**

Directory deletion cleanup at background failed if there were huge number of empty directories (millions of empty directories), whose combined metadata size crossed the Ratis allowed request size for deletion. This issue is now resolved.

**CDPD-74630: Backport HADOOP-18845 "Add ability to configure ConnectionTTL of http connections while creating S3 Client." to 7.1.9.1000****CDPD-44185: Improved session management for Hue logins with Impala**

When logging into the Hue web UI, users may occasionally encounter 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 arises when a Hue user's request is routed to a different Hue backend 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.

**CDPD-74624: IMPALA-13262 Incorrect results due to predicate pushdown in join condition**

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.

**CDPD-74309: IMPALA-13203 Boolean literals in OR conditions are not simplified as expected**

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.

**CDPD-74842: IMPALA-13302 Conjunct registration issue during query rewrite in Impala**

This issue, introduced in version 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.

**CDPD-74560: IMPALA-13378 Impalad crash due to inconsistent tuple IDs during query execution**

Avoid impalad crash in `RowDescriptor::InitTupleIdxMap()` by verifying tuple ids in descriptor table received in executor side.

**CDPD-72522: IMPALA-12582 Executor crash during runtime filter generation**

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 was resolved by ensuring ScalarExprEvaluator properly calls the Open() function, preventing out-of-bounds access and stabilizing query execution.

**CDPD-74492: [7.1.9 SP1 CHF3] [UI] 'Exclude SubTypes' and 'Exclude Sub-classifications' filter should be removed from referred attributes tab of entity details page**

1)Added the validation for glossary, term and category name. The special characters will not be allow on the glossary, term and category name field. 2)The two unused filters 'Exclude SubTypes' and 'Exclude Sub-classifications' from referred attributes tab of entity detail page were removed

**CDPD-74472: When migration restarts it results into deletion of edges and vertices**

fixed migration flow for : When migration is rerun after interruption it will avoid deleting edges and vertices from existing data and process data faster

**CDPD-74441: [UI] 'Exclude SubTypes' and 'Exclude Sub-classifications' filter should be removed from referred attributes tab of entity details page**

1)Added the validation for glossary, term and category name. The special characters will not be allow on the glossary, term and category name field. 2)The two unused filters 'Exclude SubTypes' and 'Exclude Sub-classifications' from referred attributes tab of entity detail page were removed

**CDPD-74416: Backporting rel2j - CDPD - 74200 into 7.1.9 SP1 CHF3**

When rel2j library was bumped up to higher version, Recon was failing to start. So this Recon start was fixed as part of this JIRA.

**CDPD-74377: Replace solr.version with solr\_lkgb\_jar\_version in the Ozone service.**

This issue is now resolved. The build components file now replace pom solr.version with solr\_lkgb\_jar\_version enabling the Ozone to use the downstream version of Solr

**CDPD-74354: HIVE-28484 SharedWorkOptimizer sends DPP events to unknown operators**

This issue is now fixed.

**CDPD-74344: Ozone Recon Heatmap - Make SolrHttpClient to work wirh Solrj lib when kerberos is enabled.**

When Recon Heatmap started using downstream version of solrJ lib, then kerberos client auth stopped working, This was fixed as part of this JIRA.

**CDPD-74311: 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-73882: Alter Rename should not check for the CREATE permission on the database in which renamed table is created.**

Alter rename command doesn't require CREATE permission on the database in which renamed table is created

**CDPD-73264: Backport OMID-262: Remove commons-configuration**

Upgraded Commons-configuration2 to 2.10.1.

**CDPD-72032: Searching with the application ID and DAG ID not working in Hive Queries**

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-72019: Remove the locks from SnapshotPurge and SnapshotSetProperty APIs**

This issue is now resolved by fixing Ozone Manager crash issue when the follower is lagging and it executes purgeKeys or snapshotMoveDeletedKeys for the Snapshot in one transaction.

**CDPD-71580: workaround needed for bootbox due to CVE-2023-46998**

Bootbox.js library is outdated and vulnerable. so we remove it and use new library Bootprompt.

**CDPD-69970: Support custom delimiter in SkippingTextInputFormat**

Support custom delimiter in SkippingTextInputFormat

**CDPD-68793: Hadoop - Upgrade Kafka Clients due to CVEs**

Fixes CVE-2023-25194, CVE-2021-38153 and CVE-2018-17196

**CDPD-68676: The `getTopicContent` does not always return messages when available**

Fixed an issue when `getTopicContent` didn't return all messages up to the specified end offset when an individual poll request took too much time. The timeout for the whole `getTopicContent` request defined in `responseTimeoutInMs` still applies.

**CDPD-67096: Upgrade commons-compress to 1.26.0 due to CVE-2024-25710 and CVE-2024-26308**

This issue is now resolved by fixing CVE-2024-25710 and CVE-2024-26308.

**CDPD-66299: Set role command is not audited by ranger**

Audit Set role command in ranger

**CDPD-65034: Receiving Error " TypeError: 'NoneType' object is not callable" in `TCLIService.py` when custom headers are being set**

When XSRF (Cross-Site Request Forgery) or CSRF (Cross-Site Request Forgery) is enabled in Hive or Impala, you might encounter the error "Error " TypeError: 'NoneType' object is not callable" in `TCLIService.py`. You can resolve this issue by upgrading to 7.1.9 SP1 CHF3 version.

**CDPD-64286: Hue - Metrics cannot restore database connections on failure**

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-62285: S3A StagingCommitter does not clean up staging-uploads directory**

Fixed an issue that caused incomplete cleanup of staging committer limiting the number of jobs that could be committed.

For more information, see [HADOOP-18793](#).

**CDPD-60647: Snapshot purge should be an atomic operation.**

This issue is resolved by fixing OM crash issue when the follower is lagging and it executes `purgeKeys` or `snapshotMoveDeletedKeys` for the Snapshot in one transaction.

**CDPD-8041: Add HTTP Cache-Control Header**

Added two Zeppelin configuration properties to control the HTTP response headers, allowing for improved cache management and server response behavior. See [Configuring Zeppelin caching](#).

**KUDU-3619: 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 CDP 7.1.1 and newer versions. For more information, see [KUDU-3619](#).

There are no Common Vulnerabilities and Exposures (CVE) fixed in this CHF.

**Known issues in 7.1.9 SP1 CHF 3**

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.1.9 SP1 CHF 3.

The list of known issues for the Runtime release 7.1.9 SP1 CHF 3 includes the known issues from Runtime release 7.1.9 SP1. For more information, see [Known issues in Cloudera Runtime 7.1.9 SP1](#).

**CDPD-45370: Cartesian join fails with an empty table**

When the configuration `hive.tez.cartesian-product.enabled` is set to true, query fails if one of the tables involved in a Cartesian join has zero rows. This results in an `IllegalArgumentException` with the message `Number of items is 0. Should be positive`.

To address this issue, set `hive.tez.cartesian-product.enabled` to false before running the query.

Apache Jira: [\[HIVE-26582\]](#)

**CDPD-75742: Upgrade commons-lang3 to 3.13.0**

The Commons-text 1.11.0 in Spark calls Range.of(..) method of commons-lang3 which is not available in the version 3.12.0, resulting in Oozie-Spark job failure.

Upgrade to commons-lang3 version 3.13.0 which contains the required method.

**CDPD-70450: Impala SQL queries that include the “WITH” clause should populate lineage in Atlas**

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.

**COMPX-18909: YARN NodeManagers potentially getting into unhealthy state caused by YARN-11709**

A logic was introduced to mark the NodeManager unhealthy if the container-executor encounters an unrecoverable or configuration error. For example, if the container-executor script is missing from the configured place. There is an edge-case in this behaviour that, if the application is killed just before one of the containers is trying to access the localizer syslog file, an IOException occurs because the NodeManager has already deleted the application log directory. This IOException is then caught by the logic of CDPD-49702/YARN-11709, and the NodeManager is marked unhealthy incorrectly.

Restart the unhealthy NodeManager.

**CDPD-77399: HBase fails to register the servlet metrics and throws ClassNotFoundException: org.apache.hadoop.metrics.MetricsServlet**

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](#)

**CDPD-76035: Resource lookup for Atlas service is failing**

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

**CDPD-75422: Impala schema case sensitivity issue**

Impala's schema is case insensitive, causing errors with mixed case schema elements created through Spark during predicate pushdown.

- Create tables through Impala to ensure lower case schema.
- Avoid upper case in Spark: Do not use upper case when creating tables through Spark.
- Fix existing tables: Use ALTER TABLE to rename upper case columns:

```
ALTER TABLE `database`.`iceberg_table` CHANGE COLUMN ID id string;
```

Apache Jira: [IMPALA-13463](#)

**CDPD-71411: Atlas Hbase import fails with jackson-databind exception**

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-79160: NPE while deleting BusinessMetadata**

If business metadata is created without adding any applicable types, a NullPointerException is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

**CDPD-61527: On FIPS cluster, Knox-Impala connection is failing with SSL error code 5**

The Knox-Impala connection is failing with SSL error code 5 on FIPS clusters. This prevents accessing Impala through Knox.

**CDPD-78656: Health test for Knox fails if the gateway.client.auth.needed = true is set**

The health test for Knox Gateway fails if the gateway.client.auth.needed parameter is set to true. Environments using the "curl" call are impacted. The curl call from CM is not specifying any certificate (store) while Knox is configured to require one.

To avoid the startup failure, use gateway.client.auth.wanted = true.

**CDPD-77911: Missing Log4j Redactor dependency**

The missing org.cloudera.logredactor dependency can result impossible Log4j errors during startup. Additionally, it can lead to unwanted data leak because of incorrect log redaction.

**CDPD-80921: Without permission for one glossary, /glossary call throws exception**

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-82056: UI: when server response date fields as '0', UI shows as current time**

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](#)

**CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export**

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**

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](#)

**CDPD-66938: [Analyze] [Atlas] test\_time\_range tests fail**

When the Apache Atlas server is running on a node which has time zone other than UTC, there might be a time of day when the search results might differ if the relative CreateTime date range filters of TODAY, YESTERDAY, etc. are used. 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.

Use explicit date range filters instead of using relative date range filters, such as, TODAY, YESTERDAY.

**CDPD-84502: Advanced Search is not working properly in Atlas UI**

The Advance Search in Atlas does not return any results when selecting the Typedef option. The API is not triggered for each relation.

None

**CDPD-76789: Creating tag with name description throws java.lang.ClassCastException**

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails**

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

#### **CDPD-70321: Atlas Parallel import is failing with various errors**

During a parallel import-export activity with six iceberg table policies with exportOption as db1.\* for all six exports, all import fail after the exports.

#### **CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen**

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

#### **Crash in Navigator Encrypt kernel module if --profile-file= option is used with ACLs**

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

#### **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.

#### **UnsupportedClassVersionError**

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.

#### **CDPD-87670: [Atlas UI] Apache Atlas Glossary becomes unresponsive when the page size is set to 50**

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn**

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

**CDPD-79099: Unable to import hive tables using import script in Public Cloud**

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-88230: Bulk import operations currently do not construct lineage as expected**

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**

Bulk import operations do not create audits for the imported entities.

**Technical Service Bulletins****TSB 2025-820: Potential Data Integrity Issues Found in Ozone**

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

**Knowledge article**

For the latest update on this issue see the corresponding Knowledge 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**

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.

## Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

## Cumulative hotfix 2

Know more about the cumulative hotfix 2 for 7.1.9 SP1. This cumulative hotfix was released on 30 Sep, 2024.

### Cloudera Runtime 7.1.9.1010 (Cumulative Hotfix 2) download URL:

#### Parcel Repository Location

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1010/parcels/
```

### Fixed issues in 7.1.9 SP1 CHF 2

Know more about the cumulative hotfix 2 for 7.1.9 SP1. This cumulative hotfix was released on 30 Sep, 2024.



**Note:** Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p1010.57518810.

#### CDPD-64979: CORS is too permissive for the public APIs

This issue has been 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.

#### COMPX-17702: Backport - YARN-10345 - HsWebServices containerlogs does not honor ACLs for completed jobs

Previously, the following rest APIs did not have any authorization:

- `/ws/v1/history/containerlogs/{containerid}/{filename}`
- `/ws/v1/history/containers/{containerid}/logs`



This is now resolvedx and the APIs now have an ACL authorization.

**CDPD-71849: Backport HIVE-28325 to CDH-7.1.9.1000**

A HivePrivilegeObject was created in the SemanticAnalyzer that was used for Ranger calls. Ranger used the owner as a filter when searching for objects. When the owner was missing, Ranger calls slowed down, causing a slowdown in compilation time. This issue is now resolved.

**CDPD-73943: HIVE-27994: Optimized renaming of partitioned tables**

Previously, renaming partitioned tables in Hive required fetching, deleting, and re-inserting every row in PART\_COL\_STATS. This is now optimized to improve performance, especially for tables with large datasets.

**CDPD-73942: HIVE-25225: Fix for NPE in update column stats with disabled direct SQL**

In batch loading scenarios, overhead during partition column stat updates is now reduced.

**CDPD-73941: HIVE-24663: Optimized partition column stats update for large batches**

When large number of partitions (>20K) were processed, the ColStatsProcessor ran into DB issues. This issue is now resolved by introducing smaller batches for stats gathering, preventing bulk update issues.

**CDPD-73795: Backport CALCITE-6530 to 7.1.9 SP1 CHF2**

Http Sessions created by the Avatica server never expired causing Out of Memory issues. This issue is now resolved.

**CDPD-73761: [7.1.9 SP1 CHF2] Backport KAFKA-15391 Delete topic may lead to directory offline**

Fixed an issue where deleting a topic caused the entire log directory to go offline.

**CDPD-73678: Atlas - [7.1.9 SP1 CHFx]Upgrade Spring Framework to 6.1.12/6.0.23/5.3.39 due to CVE-2024-38808 and CVE-2024-38809**

Upgraded the Spring Framework version to 6.1.12 due to CVE-2024-38808 and CVE-2024-38809.

**CDPD-73643: Filter out unused CM\_USER parameter**

The unused CM\_USER field is now removed from the /cm-configs internal endpoint.

**CDPD-73442: IMPALA-13313 Resolution of potential deadlock**

When idle\_query\_timeout was set in a session, new queries stopped responding for that session and failed. This deadlock issue occurred in long-running sessions and is now resolved.

**CDPD-73423: Ranger - Upgrade Spring Framework to 6.1.12/6.0.23/5.3.39 due to CVE-2024-38808 and CVE-2024-38809**

Upgraded the Spring-framework version to 5.3.39 due to CVE-2024-38808 and CVE-2024-38809.

**CDPD-73326: Reduce memory needed to create Ranger policy engine**

An issue led to the creation of multiple RangerResourceMatchers with identical resource specification. This issue is now resolved and the creation of multiple RangerResourceMatcher objects is now avoided by maintaining a cache of them in the RangerPluginContext object associated with the Ranger policy engine, thereby reducing policy engine's memory needs.

**CDPD-73282: Backport CALCITE-6530 HTTP Sessions are never expired in Avatica server**

The http sessions created by the Avatica server did not expire and this caused the Avatica server to run out of memory. This issue is now resolved.

**CDPD-73226: Add /tmp noexec support for Zstd and Snappy compression**

Zstandard and Snappy compression now support /tmp mounted as noexec.

**CDPD-73147: [Ranger React UI] Admin audits for "Import Delete" operation type do not display service name field**

In the Ranger React UI, in admin audits, the Service name field was missing for the audits of operation type Import Delete. This issue is now resolved and the Import delete policy logs now display the service name.



**CDPD-73144: Enhance 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, such as, database/table/column, many trie instances are consulted to retrieve policies/tags/zones associated with the resource.

Such multi-trie retrieval is optimized with a 2-pass traversal - first pass to get count and the second pass to get the actual objects. Ththerefore, the trie data structure used in Ranger policy engine is now updated to support processing of evaluators during traversal.

**CDPD-72812: SSTFilteringService was disabled by default causing snapshot chain corruption.**

SSTFilteringService was disabled by default causing snapshot chain corruption. The Ozone snapshots SSTFilteringService is now enabled by default.

**CDPD-72766: CLONE - /api/atlas/admin/version works only once after atlas restart**

API, Atlas, Admin, Version now works without Atlas restart.

**CDPD-72621: Support for default constraints in HWC table writes**

Added functionality in Hive Warehouse Connector to support default constraints during table writes, enhancing data integrity and management.

**CDPD-72596: IMPALA-12921 IMPALA-12985: Support for running Impala with locally built Ranger**

The fix adds support for locally built Ranger and for using the new constructor when instantiating RangerAccessRequestImpl.

**CDPD-72292: [Private Cloud Releases] Upgrade RequireJS due to CVE-2024-38998 and CVE-2024-38999**

Upgraded the RequireJS version due to CVE-2024-38998 and CVE-2024-38999.

**CDPD-71764: XSS vulnerability in Zeppelin : Unsanitized HTML in Markdown Paragraphs**

To enhance security, Zeppelin now integrates HTML sanitization using JSoup within the markdown interpreter. This ensures that any HTML embedded in markdown is sanitized according to a configurable blacklist.

**CDPD-71279: Proposal to Upgrade All React.js Dependent Libraries**

Upgraded the React.JS related libraries.

**CDPD-71063: Hue - Upgrade eventlet to 0.35.2 and dnspython to 2.6.1 due to CVE-2023-29483**

Security fix for CVE-2023-29483.

**CDPD-69767: Backport fix to address CVE-2023-50291**

This ticket backports the fix for CVE-2023-50291 from upstream Solr repository.

**CDPD-69634, CDPD-69412: [719 SP1 CHFx] DB constraint violation error for grant and revoke command execution**

Multiple column revoke resulted in incorrect number of columns in the generated policies, and the revoke statement only revoked if there was only one column. This issue is now resolved and multiple column revoke now generates correct results.

**CDPD-69411: [719 SP1 CHFx] Impala's authorization-related tests failed due to PSQLErrorException due to Ranger's backend database**

Impala's authorization-related tests failed due to PSQLErrorException caused by Ranger's backend database. This issue is now resolved.

**CDPD-67913: IMPALA-12554 Consolidation of Ranger policies for GRANT statements in Impala**

Impala now creates a single Ranger policy for the GRANT statement, even when multiple columns are specified, reducing the overall number of policies on the Ranger server.

**CDPD-67552: [7.1.9] Add optional wait time for replication-records-lag calculation**

Streams Replication Manager (SRM) added a new metric (replication-records-lag) about replication flows. The calculation of this metric has an impact on the performance of the SRM, so the following 3 properties were introduced to be this feature configurable. These properties can be specified in the Cloudera Manager, Configuration tab of the Streams Replication Manager service, where new lines has to be added to the "streams.replication.manager.config" property, where key-value pairs can be defined. The "replication.records.lag.calc.enabled" flag helps to turn on or off the new feature. The calculation of this metric is enabled by default. The "replication.records.lag.calc.period.ms" provides an opportunity to lower the frequency of the calculation of replication-records-lag. The default value is 0, so the metric will be provided every time, but for example with 15 sec or with longer periods, the calculation will not have a huge effect on the SRM replication latency, but the new metric will be available, so it will be not entirely switched off. The "replication.records.lag.end.offset.timeout.ms" is a way to specify the Kafka end offset timeout value in case of replication-records-lag calculation. The default value is 1 min, but lower value provides the opportunity to reduce latency, beside the risk that replication-records-lag calculation may fail. On the other hand, the higher value helps to avoid the failure of the metric calculation, but it may have bigger effect on the SRM replication latency.

**CDPD-67460: Container Balancer must move only containers with size greater than 0 bytes.**

Introduced a check on the size of the containers allowed to leave the source node during the balancing process. Added a unit test `TestContainerBalancerTask#balancerShouldMoveOnlyPositiveSizeContainers` to ensure that a balancer selects only the positive size containers. This is tested using the existing parameter `TestContainerBalancerTask#balancerShouldObeyMaxSizeLeavingSourceLimit`.

**CDPD-67113: [7.1.9] Backport KAFKA-13988: Mirrormaker 2 auto.offset.reset=latest not working**

Streams Replication Manager (SRM) did not respond to `auto.offset.reset=latest` config. This issue is now resolved.

**CDPD-66886: Solr - Backport fix for CVE-2023-50298, CVE-2023-50292, CVE-2023-50386 and CVE-2023-50291**

This ticket backports the fix for CVE-2023-50298, CVE-2023-50292, CVE-2023-50386 and CVE-2023-50291 from upstream Solr repository.

**CDPD-64939: [719 SP1 CHFx] RANGER-4585 Support multiple columns policy creation in ranger for Grant / Revoke request**

Upgrading from Cloudera Runtime 7.1.7 Service Pack 3 to 7.1.9 Service Pack 1, for a GRANT statement involving multiple columns in a table, Impala service used to create one Ranger policy for each column. This issue is now resolved.

**CDPD-63596: There are a few compile-time dependencies on spotbugs-annotation, which is LGPL and not allowed by Apache Software Foundation under [Category X](#).**

The spotbugs-annotation, an LGPL thirdparty dependency from the Ozone package is now removed.

**CDPD-62620: OM checkpoint request pauses all the background workers and then creates a checkpoint after it acquires all the necessary locks. There is a deadlock between OMDBCheckpointServlet#getCheckpoint request and RocksDBCheckpoint#pruneSstFiles background process causing a halt in the OM and no new read/write requests are taken after some time.**

This issue is now resolved. The double locking is removed and now there is no deadlock between `OMDBCheckpointServlet#getCheckpoint` request and `RocksDBCheckpoint#pruneSstFiles` background process.

**CDPD-77283: SRM can skip records when replication-records-lag endoffset fetching fails**

SRM no longer skips records during data replication when end offset fetching fails.

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

- CVE-2022-31129 - Moment JS
- CVE-2023-45857 - Axios

- CVE-2024-38999 - Require JS
- CVE-2024-38998 - Require JS
- CVE-2023-29483 - Eventlet
- CVE-2024-22243 - Spring Framework
- CVE-2024-22259 - Spring Framework
- CVE-2024-22262 - Spring Framework

### Known issues in 7.1.9 SP1 CHF 2

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.1.9 SP1 CHF 2.

The list of known issues for the Runtime release 7.1.9 SP1 CHF 2 includes the known issues from Runtime release 7.1.9 SP1. For more information, see [Known issues in Cloudera Runtime 7.1.9 SP1](#).

#### CDPD-75742: Upgrade commons-lang3 to 3.13.0

The Commons-text 1.11.0 in Spark calls Range.of(..) method of commons-lang3 which is not available in the version 3.12.0, resulting in Oozie-Spark job failure.

Upgrade to commons-lang3 version 3.13.0 which contains the required method.

#### CDPD-70450: Impala SQL queries that include the “WITH” clause should populate lineage in Atlas

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-77399: HBase fails to register the servlet metrics and throws ClassNotFoundException: org.apache.hadoop.metrics.MetricsServlet

The MetricsServlet class is a Hadoop 2-based metric servlet unavailable in Hadoop 3 deployments.

Workaround: Ignore this WARN log message during HBase Master and RegionServer startup.

Apache Issue: [HBASE-28315](#)

#### CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

#### CDPD-75422: Impala schema case sensitivity issue

Impala's schema is case insensitive, causing errors with mixed case schema elements created through Spark during predicate pushdown.

- Create tables through Impala to ensure lower case schema.
- Avoid upper case in Spark: Do not use upper case when creating tables through Spark.
- Fix existing tables: Use ALTER TABLE to rename upper case columns:

```
ALTER TABLE `database`.`iceberg_table` CHANGE COLUMN ID id string;
```

Apache Jira: [IMPALA-13463](#)

#### CDPD-71411: Atlas Hbase import fails with jackson-databind exception

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-79160: NPE while deleting BusinessMetadata

If business metadata is created without adding any applicable types, a NullPointerException is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

### On FIPS cluster, Knox-Impala connection is failing with SSL error code 5

The Knox-Impala connection is failing with SSL error code 5 on FIPS clusters. This prevents accessing Impala through Knox.

### CDPD-78656: Health test for Knox fails if the gateway.client.auth.needed = true is set

The health test for Knox Gateway fails if the gateway.client.auth.needed parameter is set to true. Environments using the "curl" call are impacted. The curl call from CM is not specifying any certificate (store) while Knox is configured to require one.

To avoid the startup failure, use gateway.client.auth.wanted = true.

### CDPD-77911: Missing Log4j Redactor dependency

The missing org.cloudera.logredactor dependency can result impossible Log4j errors during startup. Additionally, it can lead to unwanted data leak because of incorrect log redaction.

### CDPD-80921: Without permission for one glossary, /glossary call throws exception

When you don't have the permission for a glossary, the /glossary call results in a 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-82056: UI: when server response date fields as '0', UI shows as current time

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](#)

### CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export

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

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](#)

### CDPD-66938: [Analyze] [Atlas] test\_time\_range tests fail

When the Apache Atlas server is running on a node which has time zone other than UTC, there might be a time of day when the search results might differ if the relative CreateTime date range filters of TODAY, YESTERDAY, etc. are used. 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.

Use explicit date range filters instead of using relative date range filters, such as, TODAY, YESTERDAY.

### CDPD-84502: Advanced Search is not working properly in Atlas UI

The Advance Search in Atlas does not return any results when selecting the Typedef option. The API is not triggered for each relation.

None

**CDPD-76789: Creating tag with name description throws java.lang.ClassCastException**

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails**

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

**CDPD-70321: Atlas Parallel import is failing with various errors**

During a parallel import-export activity with six iceberg table policies with exportOption as db1.\* for all six exports, all import fail after the exports.

**CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen**

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

**Crash in Navigator Encrypt kernel module if --profile-file= option is used with ACLs**

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

**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.

**UnsupportedClassVersionError**

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.

**CDPD-87670: [Atlas UI] Apache Atlas Glossary becomes unresponsive when the page size is set to 50**

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn**

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

**CDPD-79099: Unable to import hive tables using import script in Public Cloud**

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-88230: Bulk import operations currently do not construct lineage as expected**

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**

Bulk import operations do not create audits for the imported entities.

**Technical Service Bulletins****TSB 2025-820: Potential Data Integrity Issues Found in Ozone**

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

**Knowledge article**

For the latest update on this issue see the corresponding Knowledge 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**

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.

## Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

## Cumulative hotfix 1

Know more about the cumulative hotfix 1 for 7.1.9 SP1. This cumulative hotfix was released on 27 Aug, 2024.

### Cloudera Runtime 7.1.9.1005 (Cumulative Hotfix 1) download URL:

#### Parcel Repository Location

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.1005/parcels/
```

### What's new in Cloudera Runtime 7.1.9 SP1 CHF1

Understand the functionalities and improvements to features of components in Cloudera Runtime 7.1.9 SP1 CHF1.

## Oozie

### 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 [Updating column types in Oozie Database](#).

## Hive

### Optimized Authorization for Adding Partitions in HiveMetaStore

You can now use the `add_partitions()` API in HiveMetaStore to efficiently authorize only new partitions, avoiding unnecessary checks for existing partitions. This improvement enhances performance and reduces authorization overhead.

### Fixed issues in 7.1.9 SP1 CHF 1

Know more about the cumulative hotfix 1 for 7.1.9 SP1. This cumulative hotfix was released on 27 Aug, 2024.



**Note:** Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p1005.56443852.

**COMPX-16285: Backport YARN-6523**

Optimized the system credentials sent in node heartbeat responses.

**CDPD-72911: Backport HDDS-9943**

Fixed potential memory leak related to token renewal.

**CDPD-72777: It's not possible to search a Hue document with the intention to copy it**

The copy action was not provided in the top search or the left assist filter. This issue is now resolved.

**CDPD-72775: Copies of shared documents are not visible to the user**

Fixed an issue where copies of shared documents were not visible to the users.

**CDPD-72731: Backport HIVE-25622**

Moved the getURIForAuth method into HiveStorageHandler and changed its signature to accept the Hive Metastore (HMS) table object instead, as it provides implementations with more flexibility around constructing the URIs.

**CDPD-72703: Altering a Kudu table with per-range hash partitions might make the table unusable**

Fixed an issue where altering a table with per-range hash bucketing by dropping or adding a particular number of columns made the table inaccessible for Kudu client applications.

**CDPD-72699: Backport CDPD-70293 to 7.1.9 SP1**

Impala executor stopped responding whenever it was restarted. This occurred when extra loads were added to the Impala cluster. This issue is now resolved.

**CDPD-72651: Backport CDPD-70019 to 7.1.9 SP1**

A shell script is now added to allow users to run authorization-related tests against specified Impala and Ranger patches.

**CDPD-72647: Backport IMPALA-13130 to 7.1.9 SP1**

Under heavy load, Impala reached maximum memory for data stream operations and DataStreamService did not differentiate between types of requests and rejected requests that helped in reducing the load. This issue is now resolved and Impala now prioritizes data stream operations.

**CDPD-72628: SMM UI integration tests are unstable locally**

There was flakiness when running the Streams Messaging Manager (SMM) UI integration tests locally. This issue is now resolved.

**CDPD-72601: Backport IMPALA-13102 to 7.1.9 SP1**

Impala could not load tables with invalid stats due to this, DROP STATS or DROP TABLE could not be performed on the table. This issue is now resolved.

**CDPD-72600: Backport IMPALA-13180 to 7.1.9 SP1**

Upgraded the PostgreSQL version to 42.5.5 due to CVE-2024-1597.

**CDPD-72595: Backport IMPALA-13107 to 7.1.9 SP1**

Impala executors failed due to invalid TExecPlanFragmentInfo received by executors with instance number equal to 0. This issue is now resolved.

**CDPD-72593: Backport IMPALA-13170, IMPALA-9441 to 7.1.9 SP1**

An InconsistentMetadataFetchException error occurred when running show databases in Impala while simultaneously executing drop database to drop the newly created database in Hive. This issue is now resolved.

**CDPD-72545: Backport IMPALA-12370 to 7.1.9 SP1 CHF1**

An option is now added to customize timezone when working with UNIXTIME\_MICROS columns of Kudu tables.



**CDPD-72544: Expand security related headers set by SMM**

The following security related headers are now added to SMM UI endpoints:

- Referrer-Policy
- Cross-Origin-Embedder-Policy
- Cross-Origin-Opener-Policy
- Cross-Origin-Resource-Policy

**CDPD-72543: Security headers are not set for static files in SMM**

Fixed SMM failure and some security related headers such as, Content-Security-Policy, X-XSS-PROTECTION, X-Content-Type-Options, X-Frame-Options and Strict-Transport-Security are now applied to static files.

**CDPD-72500: 7.1.9 SP1 CHF hue failure**

Fixed Hue failure on the CDH-7.1.9.1000 branch on CentOS7 and SLES12 platforms.

**CDPD-72481: Ranger - Upgrade Graal-SDK to 22.3.5 due to CVE-2024-20918, CVE-2024-20926 and CVE-2024-20932**

Upgraded the Graal-SDK version to 22.3.5 due to CVE-2024-20918, CVE-2024-20926 and CVE-2024-20932.

**CDPD-72388: [AUTOSYNC] Quota count can go wrong when double buffer flush takes time**

Fixed an issue that caused Ozone Quota usage to be incorrectly decremented (below actual quota usage) in a situation where the flush of Ozone Manager's DB was delayed due to heavy load.

**CDPD-72230: [AUTOSYNC] Worker queue of FullTableCache cleanup consume much memory**

An optimization is now introduced to reduce memory utilization when the Ozone Manager is under excessive load.

**CDPD-72207: ll\_service\_id is empty for an invalid notification type**

Fixed the query to fetch the latestInvalidNotificationId even when ll\_service\_id is empty. This ensures that the NameNode gets the appropriate delta's mappings.

**CDPD-72142: [AUTOSYNC] Keys from DeletedTable and DeletedDirTable of AOS should be deleted on batch operation while creating a snapshot**

On snapshot creation, the DeletedTable and DeletedDirTable of the Active Object Store (AOS) were cleared. This operation was not done in the same transaction as Snapshot create which caused orphan block objects in case of bootstrapping and lagging follower. This is now resolved and Snapshot creation and clearing of the DeletedTable and DeletedDirTable on AOS are now a single batch operation.

**CDPD-72020: [AUTOSYNC] Some containers affected by HDDS-8129 may still be in the DELETING state incorrectly**

Some containers reached the DELETING state incorrectly because of HDDS-8129. Due to this, the Storage Container Manager was not able to maintain their redundancy. This issue is now resolved and these containers are now moved back to the CLOSED state so that the Storage Container Manager can handle them correctly.

**CDPD-72008: SMM UI - Upgrade Node.js to 22.4.1/20.15.1/18.20.4 due to multiple CVEs**

Upgraded the Node.js version in Streams Messaging Manager UI to 20.15.1 due CVE-2024-27980, CVE-2024-22020, CVE-2024-36137, CVE-2024-22018 and CVE-2024-37372.

**CDPD-71994: Including search within document as we used to have in CDH**

In CDP, there was a unified search at the top, and the results did not display 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-71974: [7.1.x] Spark - OpenCSVSerde treats blank value as null**

The missing values in OpenCSVSerde are now considered as empty string.

**CDPD-71969: [snapshot] Snapshot create requests failing with Ozone Manager failover error in a system with 30000 snapshots**

The snapshot directory could not be created inside the double buffer flush thread due to a deadlock between the DoubleBuffer and the StateMachine. Ozone Manager must be restarted to get out of the deadlock.

**CDPD-71868: [Upgrade] Recon Fails to start during Ozone Start**

Recon failed to start after upgrade from 7.1.7 SP3 to 7.1.9 SP1 due to Ozone Manager DB layout upgrade. This issue is now resolved and Recon fetches a full snapshot from the Ozone Manager.

**CDPD-71807, CDPD-69782: [719 SP1 CHF1 CLONE] Users observing role change from ROLE\_SYS\_ADMIN to ROLE\_USER**

The updateUserRoleAssignments function in Ranger-admin reset the role of the user from admin to user role for users which were part of the request but were not part of the same page when paged requests were sent to Ranger-admin from Ranger-usersync. This issue is now resolved.

**CDPD-71781: [AUTOSYNC] Move SstFiltered flag to another SnapshotProperties table**

The SSTFilteringService updates the snapshotInfoTable that caused snapshot chain corruption if, Ozone Manager crashed before DB got flushed. This interfered with SnapshotInfo updated during SnapshotPurgeRequest. This issue is now resolved.

**CDPD-71738: Port CDPD-67520 to 7.1.9 SP1 CHF1**

The custom Kudu JSON Web Token (JWT) Principal claim is now usable for HBase. The hbase.security.oauth.jwt.token.principal.claim configuration property allows the use of a Subject/Principal claim different from the default sub.

**CDPD-71658: Hue QP MySQL Connector/J license issue**

The MySQL connector/J jar is not included in the GBN.

**CDPD-71433: Handle connect logical type null values in AvroConnectTranslator**

When the time.precision.mode : connect property was provided to the Debezium connector, the connect logical types were used and null values were not handled, that caused a Null Point Exception (NPE). This issue is now resolved.

**CDPD-71403: Consistent expand icon throughout the UI**

All tables in the SMM UI are now adjusted to use the same expand icon.

**CDPD-71402: Active or Inactive status does not have high contrast**

The producer/consumer status text color in the SMM UI is now adjusted to be more distinguishable.

**CDPD-71400: High contrast throughout the UI, fonts are smaller**

Adjusted font sizes and colors throughout the whole application of SMM UI to improve usability.

**CDPD-71399: Create Topic Modal issues**

Fixed visual issues in the topic creation modal.

**CDPD-71398: Create connectors should be splitted by tabs**

The connector classes in the creation wizard are now split into tabs.

**CDPD-71397: Connector metrics are misaligned**

Fixed a visual issue where connector metrics were not properly aligned.

**CDPD-71362: [7.1.9 SP1 CHF1] Temporarily disable the tasks tab on Entity Detail page**

In the **Entity Detail** page, the API of the **Tasks** tab now displays information depending on the server side property atlas.tasks.ui.tab.enabled. This property was set to false previously.

**CDPD-71340: There should be a visual clue when hitting restart on Kafka Connect**

Added a visual feedback when connectors are restarted in the form of a notification.

**CDPD-71339: Classnames on the Kafka Connect popup are not wrapped**

Previously, the connector classnames overflowed popups. This visual issue is now resolved.

**CDPD-71338: Menu items flow under some tables**

The sidebar popups were previously hidden by the main content. This visual issue is now resolved.

**CDPD-71336: Filter selector styling**

The dropdowns on the main page are now modified to make the counts more prominent.

**CDPD-71335: Listing page table styling**

Modified the table styling to for better visual.

**CDPD-71333: Wrong table headers are sticky after topic/broker dropdown**

Listing page table headers are now sticky instead of the nested table headers.

**CDPD-71332: The text in the sidepanel column headers are oversized**

Fixed a visual issue where text would overflow in the Kafka clients sidebars.

**CDPD-71117: Oozie server does not pass action start time to action conf causes a restarting launcher doesn't find child apps**

Whenever Yarn restarted the Oozie Launcher AM, Oozie could not find the previously started child jobs due to a missing original start timestamp from the Oozie Server. And the previously started child Jobs were not terminated when the Launcher AM was restarted. This issue is now resolved.

**CDPD-71062: The search component does not work with a regexp**

Added a checkbox to the listing page searchbar that toggles the regex search.

**CDPD-70955: Ozone - Upgrade Axios version to 0.28.1/1.7.2 due to CVE-2023-45857**

Upgraded the Axios version to 1.7.2 due to CVE-2023-45857.

**CDPD-70838: Impala session is hanging during catalog and statestore HA testing**

A query session on Impala did not respond when the catalogd service failed due to a standby instance if CatalogD HA was enabled. This issue is now resolved.

**CDPD-70456: Ozone Recon - Solr Health Check API throws error in upgrade from 7.1.7 to 7.1.9 version**

During an upgrade from 7.1.7 to 7.1.9 version, Solr service flag was not enabled by default even when the Solr server was running and the Solr service run was successful. Due to this Solr provider impl class failed to initialize and load, and the Solr health check API did not check if the provider was null, leading to a NPE. The Recon Overview UI did not load due to the NPE. This issue is now resolved by adding a null check.

**CDPD-70422: Cannot enforce Oozie parameter oozie.http.hostname**

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-70409: Recon Overview Page UI fails to load if Recon Solr Health throws error**

Fixed an issue where the Recon UI failed to load when there was an error in the Solr Health check API.

**CDPD-70357: [7.1.x] Do not call HMS to get list of pruned partitions when translated filter is empty**

Minimized the calls to Hive Metastore (HMS) layer to get the partitions list by making one call for each table irrespective of repetition.

**CDPD-69849: HPL/SQL: Backport HIVE-28253 to 7.1.9 SP1**

As a part of error handling in HPL/SQL mentioned in [Error Handling in HPL/SQL](#) there was a syntax error when trying to set a value for hplsql.onerror using the SET command. This error occurred because, as per the current Hplsql.g4 grammar file, SET command took only an identifier in which dot(.) was not allowed. Hence, SET hplsql.onerror did not work as per the grammar file. This issue is now resolved and the grammar file is changed to use a qualified identifier to accept dot(.).



**Note:** Use EXECUTE 'SET <HS2 Config>' or EXECUTE IMMEDIATE 'SET <HS2 Config>' for setting any HS2 related config value.

For example: EXECUTE 'SET hive.merge.split.update=true';

**CDPD-69848: Execute immediate 'select count from tbl' throwing ClassCastException in hplsql mode.**

There was a java.lang.ClassCastException error when SELECT count(\*) from result was executed. This is ClassCastException is caused because, the select count(\*) query returns a long value but HPLSQL expects a string type value.

As a workaround, instead of using EXECUTE IMMEDIATE 'SELECT count(\*) from result', use the query directly SELECT count(\*) from result.

**CDPD-69847: Signalling CONDITION HANDLER is not working in HPLSQL**

The CONDITION HANDLERS defined by a user were not invoked when the signal was given to the corresponding condition. As a part of processing the exception/error handling conditions, user-defined conditions are also processed (just logging into the log file). And, when it has to invoke user-defined condition handlers there are no condition handlers in the stack.

**CDPD-69846: HPL/SQL: Backport HIVE-28214 to 7.1.9 SP1**

Previously, HPLSQL did not use the Hive variables that were passed through Beeline but used the --hivevar option. This issue is now resolved by modifying the code to use the Hive variables passed through Beeline.

**CDPD-68849: Set kerberosEnableCanonicalHostnameCheck=false in beeline on CDP Base**

Previously, there was no support when using Kerberos authorisation when connecting to the Hive Virtual Warehouse (VW) using Private Cloud Base Beeline. This issue is now resolved and kerberosEnableCanonicalHostnameCheck is now set to false.

**CDPD-67602: Hue - Upgrade PostgreSQL to 42.5.5/42.6.1/42.7.2 due to CVE-2024-1597**

Upgraded the PostgreSQL version to 42.5.5/42.6.1/42.7.2 due to CVE-2024-1597.

**CDPD-66706: HDFS file compression in Hue is failing**

Hue HDFS file compression feature did not work and failed with an error message in the OOZIE application logs. This issue is now resolved.

**CDPD-66321: IMPALA-12559 Support x5c Parameter in JSON Web Keys (JWK)**

Impala coordinator failed to parse in JSON Web Key Set (JWK) and failed to start the coordinator. This issue is now resolved.

**CDPD-62360: When Ranger authorisation is disabled for Solr, we get NPE when calling Solr API using user which does not have permission**

There was an exception when calling Solr API when Ranger authorisation was disabled for Solr. This issue is now resolved.

**CDPD-61728: Upgrade node.js to fix high cve**

Upgraded the Node.js version due to CVE-2021-3450, CVE-2021-44531, CVE-2023-32004 and CVE-2023-32006.

**CDPD-57994: Hue import not creating External tables**

Previously, there was 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-57931: Upgrade OpenJPA in Oozie, because of CVE**

Upgraded the OpenJPA version to 3.2.2 due to CVE.

**CDPD-49745: Expand app\_path column in \*\_JOBS tables to allow HDFS paths longer than 255 characters**

The APP\_PATH column now supports storing paths longer than 255 characters.

**CDPD-48664: Retry mechanism anomaly in Oozie with High Availability enabled**

There was an issue with the retry mechanism in Oozie when High Availability was enabled. This issue is now resolved.

**CDPD-48112: Hue - Upgrade Bootstrap to 3.4.1 or 4.3.1+ due to CVE-2019-8331**

Upgraded the Bootstrap version 4.3.1 due to CVE-2019-8331.

**CDPD-68483: Avro: hadoop-client version conflict**

This fix resolves an issue related to the avro-tools uber JAR having two hadoop-client JARs with different versions, which was causing a conflict.

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

- CVE-2023-51775 - Jose4J
- CVE-2023-22102 - MySQL-Connector-J
- CVE-2024-1597 - PostgreSQL

**Known issues in 7.1.9 SP1 CHF 1**

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.1.9 SP1 CHF 1.

The list of known issues for the Runtime release 7.1.9 SP1 CHF 1 includes the known issues from Runtime release 7.1.9 SP1. For more information, see [Known issues in Cloudera Runtime 7.1.9 SP1](#).

**CDPD-75742: Upgrade commons-lang3 to 3.13.0**

The Commons-text 1.11.0 in Spark calls Range.of(..) method of commons-lang3 which is not available in the version 3.12.0, resulting in Oozie-Spark job failure.

Upgrade to commons-lang3 version 3.13.0 which contains the required method.

**CDPD-70450: Impala SQL queries that include the “WITH” clause should populate lineage in Atlas**

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-77399: HBase fails to register the servlet metrics and throws ClassNotFoundException: org.apache.hadoop.metrics.MetricsServlet**

The MetricsServlet class is a Hadoop 2-based metric servlet unavailable in Hadoop 3 deployments.

Workaround: Ignore this WARN log message during HBase Master and RegionServer startup.

Apache Issue: [HBASE-28315](#)

**CDPD-76035: Resource lookup for Atlas service is failing**

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

**CDPD-75422: Impala schema case sensitivity issue**

Impala's schema is case insensitive, causing errors with mixed case schema elements created through Spark during predicate pushdown.

- Create tables through Impala to ensure lower case schema.
- Avoid upper case in Spark: Do not use upper case when creating tables through Spark.
- Fix existing tables: Use ALTER TABLE to rename upper case columns:

```
ALTER TABLE `database`.`iceberg_table` CHANGE COLUMN ID id string;
```

Apache Jira: [IMPALA-13463](#)

**CDPD-71411: Atlas Hbase import fails with jackson-databind exception**

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-79160: NPE while deleting BusinessMetadata**

If business metadata is created without adding any applicable types, a NullPointerException is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

#### **On FIPS cluster, Knox-Impala connection is failing with SSL error code 5**

The Knox-Impala connection is failing with SSL error code 5 on FIPS clusters. This prevents accessing Impala through Knox.

#### **CDPD-78656: Health test for Knox fails if the gateway.client.auth.needed = true is set**

The health test for Knox Gateway fails if the gateway.client.auth.needed parameter is set to true. Environments using the "curl" call are impacted. The curl call from CM is not specifying any certificate (store) while Knox is configured to require one.

To avoid the startup failure, use gateway.client.auth.wanted = true.

#### **CDPD-77911: Missing Log4j Redactor dependency**

The missing org.cloudera.logredactor dependency can result impossible Log4j errors during startup. Additionally, it can lead to unwanted data leak because of incorrect log redaction.

#### **CDPD-80921: Without permission for one glossary, /glossary call throws exception**

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-82056: UI: when server response date fields as '0', UI shows as current time**

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](#)

#### **CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export**

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**

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](#)

#### **CDPD-66938: [Analyze] [Atlas] test\_time\_range tests fail**

When the Apache Atlas server is running on a node which has time zone other than UTC, there might be a time of day when the search results might differ if the relative CreateTime date range filters of TODAY, YESTERDAY, etc. are used. 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.

Use explicit date range filters instead of using relative date range filters, such as, TODAY, YEST ERDAY.

#### **CDPD-84502: Advanced Search is not working properly in Atlas UI**

The Advance Search in Atlas does not return any results when selecting the Typedef option. The API is not triggered for each relation.

None

#### **CDPD-76789: Creating tag with name description throws java.lang.ClassCastException**

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails**

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

#### **CDPD-70321: Atlas Parallel import is failing with various errors**

During a parallel import-export activity with six iceberg table policies with exportOption as db1.\* for all six exports, all import fail after the exports.

#### **CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen**

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

#### **UnsupportedClassVersionError**

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.

#### **Crash in Navigator Encrypt kernel module if --profile-file= option is used with ACLs**

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

#### **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**

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn**

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

**CDPD-79099: Unable to import hive tables using import script in Public Cloud**

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-88230: Bulk import operations currently do not construct lineage as expected**

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**

Bulk import operations do not create audits for the imported entities.

**Technical Service Bulletins****TSB 2025-820: Potential Data Integrity Issues Found in Ozone**

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

**Knowledge article**

For the latest update on this issue see the corresponding Knowledge 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**

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.

**Knowledge article**

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

## 7.1.9

These are Release Notes for CDP Private Cloud Base 7.1.9.

Information for CDP Private Cloud Base 7.1.9 are found here.

## What's new in Cloudera Runtime 7.1.9

Understand the functionalities and improvements to features of components in Cloudera Runtime 7.1.9.

### Open Data Lakehouse, powered by Apache Iceberg

CDP Private Cloud Base 7.1.9 delivers the hybrid Open Data Lakehouse providing the following benefits:

**Open architecture**

Cloudera's Open Data Lakehouse, powered by Apache Iceberg is 100% open—open source, open standards based, and with wide community adoption. It can store multiple data formats and enables multiple engines to work on the same data.

**Ease of adoption**

By integrating Iceberg right into the Shared Data Experience (SDX) and Apache Ozone, Cloudera offers the easiest path to deploying a lakehouse. Additional capabilities like schema evolution, hidden partition, and more simplify data management for large datasets.

**Multi-cloud**

Build a lakehouse in your own data center and run anywhere. Cloudera offers the same data services with complete portability on all clouds.

**Secure and governed**

Iceberg tables integrate within SDX, allowing for unified security, fine-grained policies, governance, lineage and metadata management across multiple clouds, so you can focus on analyzing your data while Cloudera takes care of the rest.

For more information, see [What is Open Data Lakehouse in CDP Private Cloud Base?](#)

### Apache Ozone

**Ozone Snapshots**

Ability to create, manage, and delete snapshots at Volume and Bucket levels

**Ozone Quotas**

The Ozone shell is the primary command line interface for managing the quota of volumes and buckets.

### Container Balancer Threshold

You can now determine the threshold value before configuring the required parameters

### Erasure Coding

The Ozone Erasure Coding feature provides data durability and fault-tolerance along with increased storage efficiency

### SCM and OM decommissioning

Gracefully remove an OM/SCM from the SCM HA Ring.

For more information on the Ozone 7.1.9 features, see [What's New in Ozone](#).

### Zero Downtime Upgrades (ZDU)

Cloudera is reintroducing the concept of rolling upgrades in CDP 7.1.9 in an easier to use format called Zero Downtime Upgrades (ZDU). Zero Downtime Upgrades automates the process of performing rolling upgrades in an optimized format to allow for minimal to zero downtime depending on the services installed on a cluster. All future service packs and runtime upgrades will support ZDU. However, the enhancements brought by ZDU will be available on upgrades from CDP 7.1.7 and CDP 7.1.8. Before using this feature read the upgrade instructions. For more information, see the [Zero Downtime upgrade](#) documentation.



**Caution:** Cloudera recommends all upgrades to happen in a maintenance window by throttling and scaling down workloads during that time as a best practice.

### Upgrade CDP 7.1.7 SP2 to CDP 7.1.9

You can perform an In-place upgrade from CDP 7.1.7 SP2 to CDP 7.1.9. For more information, see [CDP to CDP](#) documentation.

### Upgrade CDP 7.1.8 to CDP 7.1.9

You can perform an In-place upgrade from CDP 7.1.8 to CDP 7.1.9. For more information, see [CDP to CDP](#) documentation.

### Upgrade CDP 7.1.7 SP1 to CDP 7.1.9

You can perform an In-place upgrade from CDP 7.1.7 SP1 to CDP 7.1.9. For more information, see [CDP to CDP](#) documentation.

### Upgrade CDH 6 to CDP 7.1.9

You can perform In-place upgrade from CDH6 to CDP 7.1.9. For more information, see [CDH 6 to CDP](#) documentation.

### Upgrade HDP 3 to CDP 7.1.9

You can perform In-place one-stage upgrade from HDP 3 to CDP 7.1.9 using CMA 2.6.2. For more information, see [HDP 3 to CDP](#) documentation.

### Rollback CDP 7.1.9 to CDP 7.1.7 SP2

You can downgrade or rollback an upgrade from CDP Private Cloud Base 7 to CDP 7.1.7 SP2. The rollback restores your CDP cluster to the state it was in before the upgrade, including Kerberos and TLS/SSL configurations. For more information, see [Rollback CDP 7.1.9 to CDP 7.1.7 SP2](#) documentation.

### Rollback CDP 7.1.9 to CDP 7.1.8

You can downgrade or rollback an upgrade from CDP Private Cloud Base 7 to CDP 7.1.9. The rollback restores your CDP cluster to the state it was in before the upgrade, including the Kerberos and TLS/SSL configurations. For more information, see [Rollback CDP 7.1.9 to CDP 7.1.8](#) documentation.

### Rollback CDP 7.1.9 to CDH 6

You can downgrade or rollback an upgrade from CDP Private Cloud Base 7 to CDH 6. The rollback restores your CDH cluster to the state it was in before the upgrade, including the Kerberos and TLS/SSL configurations. For more information, see [Rollback CDP 7.1.9 to CDH 6](#) documentation.

### TLS 1.2 support for secured database connections

Certain CDP components, Cloudera Manager server, and Reports Manager support connections to backend databases that are secured with Transport Layer Security (TLS) 1.2 encryption. This enhances the security connections between CDP and backend databases, such as MySQL, PostgreSQL, MariaDB.

The following CDP components are supported:

- Hive MetaStore
- Hue
- Schema Registry
- Streams Messaging Manager
- Oozie
- Sqoop
- Ranger
- Ranger KMS

For more information on cluster services, see [Configuring TLS 1.2 for cluster services](#)

For more information on Cloudera Manager, see [Configuring TLS 1.2 for Cloudera Manager](#)

For more information on Reports Manager, see [Configuring TLS 1.2 for Reports Manager](#).

### TCPS support for connections to Oracle database

Certain CDP components support connections to backend Oracle database that are secured with Transmission Control Protocol with SSL (TCPS). This provides greater security for connections between CDP and the backend Oracle database. For more information, see [TCPS support for connections to Oracle database](#).

The following CDP components are supported:

- Hive MetaStore
- Hue
- Schema Registry
- Streams Messaging Manager
- Oozie
- Sqoop
- Ranger
- Ranger KMS

### Collecting Heartbeat data from Cloudera Manager

Beginning with Cloudera Manager 7.11.3, a report containing basic cluster information securely transmits to Cloudera periodically. This report contains cluster-related metadata to determine the version and size of each cluster. This information helps Cloudera to gain a clearer understanding of your deployments and deliver more robust support and ensure an improved customer experience.

Reports are saved locally for you with infrastructure isolated from the public internet. For assistance, open a General Administrative Assistance case on [MyCloudera](#).

The generated report is human-readable for users and can be found under `/var/lib/cloudera-scm-server/reports` (configured as default).

### Client RPMs

Client RPMs enable you to access services installed on the CDP cluster from your containerized application or edge nodes that are not managed by Cloudera Manager. You can use Client RPMs for the use cases that require thin client applications binding without having to download and install the entire CDP parcel which requires over 10 GB of storage space on each platform node. Starting with the CDP Private Cloud Base 7.1.9 release, '-client' RPM/Debian packages are available.

For more information, see [Application Access](#)

### Support for noexec option on the /tmp directory

All Cloudera Runtime services, with the exception of Oozie and Sqoop, support the noexec option on the /tmp directory. The /tmp directory on Linux hosts is used by many applications to store non-persistent data and to execute transient scripts. Users require the noexec option on /tmp directory to eliminate possible security risks by preventing the execution of binaries from the /tmp filesystem. The noexec option prevents unintentional system modifications or corruption that may potentially lead to system instability or data theft.

### Deprecation notices

For information, see [Deprecation notices in Cloudera Runtime 7.1.9](#).

## What's New in Apache Atlas

Learn about the new features of Apache Atlas in Cloudera Runtime 7.1.9.

### Storage reduction for Atlas

Audit aging reduces the existing audit data in the Atlas system which is based on the end user criteria and configuration changes that users can manage.

For more information, see [Storage reduction for Atlas](#).

### Iceberg support for Atlas

Atlas integration with Iceberg helps you identify the Iceberg tables to scan data and provide lineage support.

For more information, see [Iceberg support for Atlas](#).

### Ability to download search results from Atlas UI

Atlas supports improved search results capabilities. For more information, see [Ability to download search results](#).

### Using Relationship search

Entities in Atlas can be searched based on the relationships that describe various metadata between a couple of entity end-points.

For more information, see [Relationship search](#).

### Log4j updates

Log4j 1.x has been migrated to reload4j.

## RHEL-9.1 operating system support

Atlas is now supported on RHEL-9.1.

## Support for validating the AttributeName in parent and child TypeDef

Atlas service validates the attribute names of the entity types for those attributes having identical names as their parents' attributes.

For more information, see [Support for AttributeName in parent and child TypeDef](#).

## Custom Kerberos Principal changes for Atlas

Customizing the Atlas Kerberos Principal is a two-step process. For more information, see [Custom Kerberos Principal for Atlas](#).

## What's New in Cruise Control

Learn about the new features of Cruise Control in Cloudera Runtime 7.1.9.

### Cruise Control 2.5.116 Rebase

Cruise Control in Cloudera Runtime is rebased to the 2.5.116 version. For more information about the fixes and features in Cruise Control 2.5.116, see the [Cruise Control Rebase Summary](#).

### New endpoint for Cruise Control

The GET/kafkacruisecontrol/permissions endpoint is added to Cruise Control that lists the level permissions of the current user. In case authentication is not configured for a user, the GET call returns Unable to retrieve privilege information for an unsecure connection message.

### Adding MultiLevelRackAwareGoal to Cruise Control

As the currently available RackAwareGoal in Cruise Control does not support multi level rack awareness, a new goal was created to ensure that the replicas are assigned respecting the rules of multiple level racks.

The new goal is named com.cloudera.kafka.cruisecontrol.analyzer.goals.MultiLevelRackAwareDistributionGoal.

For more information, see the [Multi-level rack-aware distribution goal](#) documentation.

### The replication factor of Cruise Control internal topics is configurable

The following Cloudera Manager properties are introduced in the Cruise Control configuration:

- Sample Store Topic Replication Factor
- Metrics Topic Replication Factor
- Metrics Topic Partition Count
- Metrics Topic Auto Creation
- Maintenance Event Topic Replication Factor

The properties configure the replication factors of internal topics. The new default value of the replication factor is 3. If you have less than 3 Kafka brokers in your cluster, reduce the replication factor in the Cruise Control configuration. The configuration properties are automatically applied for new deployments. However, when upgrading an existing cluster, the partitions must be reassigned for the new properties to take effect.

### Cruise Control now supports Kerberos auth-to-local (ATL) rules with SPNEGO authentication

Cruise Control now uses the cluster-wide Kerberos auth-to-local (ATL) rules by default. A new configuration property called SPNEGO Principal to Local Rules is introduced. This property is used to manually specify the ATL rules. During an upgrade, the property is set to DEFAULT to ensure backward compatibility. Following an upgrade, if you want to use the cluster-wide rules, clear the existing value from the SPNEGO Principal to Local Rules property.

### Cruise Control Rebase Summary

In Cloudera Runtime 7.1.9, Cruise Control is rebased to the 2.5.116 version. Other than the added new feature, several issues are fixed and several features are enhanced to have a better performance when using Cruise Control.

**Table 4: Fixed Issues**

|                                                                                                |
|------------------------------------------------------------------------------------------------|
| <a href="#">PR-1758</a> Fix request log to write to configured CruiseControlPublicAccessLogger |
| <a href="#">PR-1847</a> Fix minor logging issues                                               |
| <a href="#">PR-1875</a> Fix the BrokerFailureDetector doesn't work issue                       |
| <a href="#">PR-1901</a> Fix failedBrokers.txt permissions                                      |
| <a href="#">PR-1939</a> Address issues with Vertx based Swagger UI                             |
| <a href="#">PR-1949</a> Fix the response code to 202 for in-progress request                   |

**Table 5: Version Update**

|                                                                                                              |
|--------------------------------------------------------------------------------------------------------------|
| <a href="#">PR-1810</a> Bump vulnerable transitive jackson-databind dependency                               |
| <a href="#">PR-1811</a> BPin jackson-databind 2.12.6.1 to resolve CVE-2020-36518                             |
| <a href="#">PR-1815</a> Bump swagger-parser for CVE-2020-36518                                               |
| <a href="#">PR-1855</a> Upgrade Netty and Jetty versions for CVE fixes                                       |
| <a href="#">PR-1861</a> Downgrade netty and jetty dependencies to a safer version                            |
| <a href="#">PR-1902</a> Bump fastxml dependencies to fix a High CVE score for the org.yaml:snakeyaml package |
| <a href="#">PR-1926</a> chore: bump fastxml dependencies to fix a High CVE score                             |
| <a href="#">PR-1928</a> Bumped dependency swagger-parser-v3 version to latest 2.1.3                          |
| <a href="#">PR-1937</a> Bump scala version to 2.13.10 (CVE-2022-36944)                                       |

**Table 6: Goal Improvements**

|                                                                          |
|--------------------------------------------------------------------------|
| <a href="#">PR-1809</a> Add BrokerSetAwareGoal                           |
| <a href="#">PR-1857</a> Add more details to the goals config description |
| <a href="#">PR-1919</a> Add fixability metrics to GoalOptimizer          |

**Table 7: Feature Maintenance**

|                                                                                                                 |                                                                                                              |
|-----------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------|
| <a href="#">PR-1789</a> Update BrokerFailureDetector to use AdminClient clusters                                | <a href="#">PR-1897</a> Add more logging to SlowBrokerFinder                                                 |
| <a href="#">PR-1803</a> Cleanups for running without ZooKeeper                                                  | <a href="#">PR-1895</a> Rename the getGoalsByPriority to getDefaultGoalsByPriority                           |
| <a href="#">PR-1825</a> Remove execution tasks that has in-movement partitions from re-execution candidate      | <a href="#">PR-1889</a> Display the brokerSet id in kafka_cluster_state endpoint                             |
| <a href="#">PR-1830</a> Make the retrieval of the desired replication factor of sample store topics more robust | <a href="#">PR-1910</a> Future-proof for Kafka 3.3 for KafkaYammerMetrics class name changes                 |
| <a href="#">PR-1831</a> Allow fraction CPU core capacity values                                                 | <a href="#">PR-1914</a> Fix leader replica cpu util when leader egress is zero                               |
| <a href="#">PR-1839</a> Update the numCore type to double in BasicStats                                         | <a href="#">PR-1921</a> Add Vertx.io based API with swagger UI                                               |
| <a href="#">PR-1841</a> Update NumCore type in brokerStats.yaml                                                 | <a href="#">PR-1941</a> Log broker-set resolution error only when BrokerSetAwareGoal is in default goal list |
| <a href="#">PR-1848</a> Enable users to get remote-storage-enabled of topics in kafka_cluster_state             | <a href="#">PR-1967</a> Add metrics to reflect the partition movement speed                                  |

|                                                                                                                            |                                                                                                           |
|----------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| <a href="#">PR-1867</a> Make Prometheus broker cpu metric query configurable                                               | <a href="#">PR-1968</a> Implement concurrency adjuster for each individual broker                         |
| <a href="#">PR-1879</a> Handle the inconsistency between clusterModel and kafka metadata during update topic configuration | <a href="#">PR-1977</a> Add brokerSet to JSON API response                                                |
| <a href="#">PR-1881</a> Mute KafkaCruiseControlConfig logs during anomaly detection                                        | <a href="#">PR-1980</a> Dynamically adjust the executionProgressCheckIntervalMs based on execution status |
| <a href="#">PR-1884</a> Make sure the provision response is always RIGHT_SIZED if it is not OVER_PROVISIONED               | <a href="#">PR-1983</a> Add initialized state to concurrency adjuster and avoid null for the metric value |
| <a href="#">PR-1891</a> Make the state output easier to read                                                               | <a href="#">PR-1985</a> Remove initialized check when get concurrency summary                             |

## What's New in Apache HBase

Learn about the new features of Apache HBase in Cloudera Runtime 7.1.9.

### 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

## What's New in Apache Hive

Learn about the new features of Hive in Cloudera Runtime 7.1.9.

### HiveServer graceful shutdown

You can now configure the graceful shutdown timeout property for HiveServer (HS2), which ensures that HS2 waits for a specified time period before shutting down, thereby allowing queries that are already running to complete before HS2 stops. For more information, see [Configuring graceful shutdown property for HiveServer](#)

### Support for HPL/SQL

Cloudera Data Platform (CDP) supports Hive Hybrid Procedural SQL (HPL/SQL). HPL/SQL is an Apache open source procedural extension for SQL for Hive users. HPL/SQL includes imperative programming structures (variables, procedures, control flow, and exceptions), and is typically used for ETL. The HPL/SQL language

understands the syntax and semantics of most procedural SQL dialects, such as Oracle PL/SQL. For more information, see [HPL/SQL stored procedures](#)

### Hive Metastore dynamic leader election

You can use Hive Metastore (HMS) leader election to avoid running the same tasks across all Hive Metastore (HMS) instances by either configuring a HMS leader manually or enabling dynamic election. Dynamic leader election feature uses Hive locks to dynamically elect a leader. When a HMS instance owns a lock, it is elected as the leader and performs the housekeeping tasks. The HMS regularly sends heartbeats to prevent the lock from timing out. For more information, see [Hive Metastore leader election](#).

### Hive audit logging improvement

You can configure the HDFS audit logs to include the caller context in the logs. This allows Hive to audit partitions that are scanned as part of a Hive query. The audit information comprises the Query ID and User ID, which helps in meeting compliance requirements, such as controlling user access to data from specific regions or access only to particular time periods. For more information, see [Configuring query audit logs to include caller context](#).

### Hive Metastore supports TLS 1.2 for secured database connections

Hive Metastore supports connections to backend databases that are secured with Transport Layer Security (TLS) 1.2 encryption. This enhances the security connections between CDP and backend databases, such as MySQL, PostgreSQL, MariaDB. For more information, see [Configuring Hive Metastore to connect to TLS-enabled database](#).

### Hive Metastore supports connections to a TCPS-enabled Oracle database

Hive Metastore supports connections to backend Oracle databases that are secured with Transmission Control Protocol with SSL (TCPS). This provides greater security for connections between CDP and the backend Oracle database. For more information, see [Configuring Hive Metastore to connect to TCPS-enabled Oracle database](#).

## What's New in Hue

Learn about the new features of Hue in Cloudera Runtime 7.1.9.

### Ability to browse Ozone filesystem from Hue

You can now browse files and directories on the Ozone filesystem from Hue, just like how you browse files on S3 or ADLS Gen2. To enable browsing files on Ozone, you must first enable the Ozone File Browser. See [Enabling browsing Ozone from Hue on CDP Private Cloud Base](#).

### Ability to create Iceberg tables using Hue Importer

You can create Iceberg tables in Hue starting with 7.1.9. You can also create an Iceberg table by importing a CSV file and selecting the Iceberg table option on the Importer. For more information, see [Creating Iceberg tables in Hue](#).



**Note:** In the CDP 7.1.9 release, Iceberg is supported only with Impala.

### Hue supports Hive Hybrid Procedural SQL

You can run Hive Hybrid Procedural SQL (HPL/SQL) using the Hue query editor. To enable the HPL/SQL interpreter, see [Enabling stored procedures for Hive on CDP Private Cloud Base](#). To run stored procedures from Hue, see [How to run a stored procedure from Hue in CDP Private Cloud Base](#).

### Hue supported on RHEL 9 with Python 3.9

Hue is now supported with Python 3.9 on RHEL 9 operating systems. RHEL 9 comes preinstalled with Python 3.9, so you do not have to install it separately. If you want to install Python 3.9 in a custom location, then see [Installing standard Python 3.9 binary on RHEL 9 at a standard or custom location](#).



### Hue supported on SLES 15

Hue is supported to run on SLES 15.

### Hue supports connections to databases secured using TLS 1.2 and TCPS

Hue can connect to TLS-enabled MySQL, MariaDB, or PostgreSQL databases and TCPS-enabled Oracle database. To connect to a TLS 1.2/TCPS-enabled database while adding the Hue service to a cluster, see [Configure TLS 1.2 for Hue](#). You can also enable TLS1.2/TCPS on an existing database and then configure Hue to connect to it. See [Set up and configure TLS 1.2 for Hue](#).

For more information about Oracle TCPS, see [How to connect CDP components to a TCPS-enabled Oracle database](#).

### Hue Query Processor supports MySQL and Oracle as backend databases

Earlier, the Hue Query Processor only used the PostgreSQL as its backend database. You can now use MySQL and Oracle as a backend databases for the Hue Query Processor on CDP Private Cloud Base. For information on how to add the Query Processor service with all the supported backend databases, see [Adding Query Processor service to a cluster](#).

### Ability to control caching behavior of the web page

You can enable web page-caching to ensure that your browser fetches latest resources while you are exploring data using Hue. Hue uses Cache-Control, Pragma, and Expires HTTP headers. To enable cache control, see [Enabling cache-control HTTP headers when using Hue](#).

### The “enable\_queries\_list” configuration has been removed from Hue jobbrowser safety valve section

The enable\_queries\_list configuration in the Hue's Advanced Configuration Snippet displayed or hid the **Queries** tab on the **Job Browser** page. This configuration has been removed. The **Queries** tab is displayed by default. You can override the query\_store configuration and hide the Queries tab. For more information, see [Hue configurations in CDP Runtime](#).

### Ability to set the file size for upload using Hue File Browser

You can set the permitted scope of a file that your users can upload using the Hue File Browser by setting the following parameter in the Hue Advanced Configuration in Cloudera Manager:

```
[filebrowser]
max_file_size_upload_limit=[***FILE-SIZE-IN-BYTES***] \\default is -1 (no
limit)
```

For more information, see [Hue configurations in CDP Runtime](#).

### Increased the download limit on the Solr dashboard

Earlier, you could download only 1000 records from the Solr Search dashboard. Hue now supports downloading up to 15000 records. You can configure the download limit using the following Advanced Configuration Snippet:

```
[search]
download_limit=[***DOWNLOAD-LIMIT***]
```

For more information, see [Hue configurations in CDP Runtime](#).

### Hue Query Processor scan frequency decreased to 5 minutes

The Hue Query Processor scans the event processor pipeline to retrieve the Hive query history and query details and displays them on the **Job Browser** page. The scan frequency has been decreased from 2 milliseconds to 5 minutes to

optimize resource utilization. As a result, you may notice a delay in viewing the query history and query details on the **Job Browser** page for queries that finish executing in less than 5 minutes. However, you can still view the query history from the **Query history** tab below the query editor. See [Configuring the Hue Query Processor scan frequency](#).

### Query Processor API to force data cleanup

Hue Query Processor cleans up queries older than a set number of days as per the set schedule. However, to manually clean up queries on an as needed basis, you can use the Query Processor API. When you call this API, it also runs a VACUUM command on the Query Processor tables. All queries that were run before the epoch time are cleared. For more information, see [Ways to clean up old queries from the Query Processor tables](#).

### Improvements and enhancements

#### SparkSQL improvements

Hue now has a dedicated autocomplete and syntax checker for Spark SQL. Hue supports all Spark SQL statement types and it is up-to-date with version 3.3.1 of the SparkSQL syntax.

The UDF library for Spark SQL has also been integrated with the autocomplete code, as well as in the right assist panel. The inline help includes all built-in functions of Spark SQL 3.3.1.

Additionally, numerous improvements have been made in the Hue backend integration with SparkSQL, including caching of session details and an overall performance boost.

#### Added an option to download documents from the Hue web interface

You can now download saved documents when you search for them from the search bar or from the left-assist panel by right-clicking on the document.

#### Keyboard shortcut for showing and hiding left and right-assist panels

You can press command + . (Macintosh) or Ctrl + . (Windows) to hide and show the left and right-assist panels. Hiding the left and right-assist panels provides a larger area for writing queries and viewing results.

## What's New in Apache Iceberg

Learn about the new feature Iceberg in Cloudera Runtime 7.1.9.

### Apache Iceberg general availability

Apache Iceberg 1.3 is now generally available (GA) in [CDP Private Cloud Base](#). Iceberg integration with CDP enhances the Cloudera open Lakehouse architecture in your own data center, or for a hybrid on-prem/cloud use case, by extending multifunction analytics to petabyte scale.

Apache Iceberg integration in CDP brings with it industry first support for:

- Apache Spark applications with full support for both read and write operations
- High-performance BI query engine powered by Apache Impala
- Streaming analytics powered by Apache Flink
- Real time streaming powered by Apache Nifi
- [Iceberg replication using Replication Manager](#)

From Impala, Spark, Flink\*, or Nifi\*, you can now use Apache Iceberg features, which include time travel, rollback, in-place table migration, schema evolution, and in-place partition evolution. Iceberg queries from Apache Hive are not supported in this release.

\* Standalone SKU required

## What's New in Apache Impala

Learn about the new features of Impala in Cloudera Runtime 7.1.9.

## Java Dependencies - JDK 17 support

Although Impala is primarily written in C++, it does use Java to communicate with various Hadoop components. From this release, the officially supported JVMs for Impala will include JDK 17 along with the existing Oracle and OpenJDK variants of 8 and 11. When adding hosts to a new cluster, you will be prompted to choose the JDK version, and you can configure any of the supported JDK versions. Internally, the `impalad` daemon relies on the `JAVA_HOME` environment variable to locate the system Java libraries.

## RHEL 9 support

From this release, Python 3-based impala clients (`Impyla` and `impala-shell`) are compatible with RHEL 9 and can be installed on RHEL 9 to communicate with an Impala instance on CDP 7.1.9.



**Note:** RHEL 9 is supported by CDP 7.1.9 and above, and RHEL 9 is compatible only with Python 3. So you may run into issues if you run `impala-shell` on RHEL 9 by building a custom Python 2, installing it yourself, and changing the Python path to Python 2. If you run into such issues, set this parameter `IMPALA_PYTHON_EXECUTABLE=python3` pointing to Python 3.

## SLES 15 support

SLES 15 is supported from CDP 7.1.8-CHF8 and above!

## Python 3.8 support on all CDP certified OSs

Note to an Impala user using `impala-shell` from the parcel: Python 3 is supported from 7.1.8 CHF3 and the CDP versions prior to this hotfix release support Python 2. So if you must use Python 3, make sure to use a version of CDP that supports it.

Note to an Impala user using `impala-shell` downloaded from <https://pypi.org/project/impala-shell/> With any version of CDP 7.x, you can use the latest PyPi version. If you must use Python 3 in your environment, make sure to use the latest PyPi `impala-shell`.

## Impala WebUI improvements

This release enhanced the Impala daemon's Web UI to display the following additional details:

- Backends start time and version: In a large cluster, you can now use the Impala daemon's Web UI to view the start time and version for all the backends.
- Query performance characteristics: For a detailed report on how a query was executed and to understand the detailed performance characteristics of a query, you can use the built-in web server's UI and look at the timeline shown in the [Gantt chart](#). This chart is an alternative to the `PROFILE` command and is a graphical display in the WebUI that renders timing information and dependencies.
- Export query plan and timeline: To understand the detailed performance characteristics for a query, you issue the `PROFILE` command in `impala-shell` immediately after executing a query. As an alternative to the profile download page, this release added support for exporting the graphical query plan and also for downloading the timeline in SVG/HTML format. Once you export the query plan or the timeline, memory resources consumed from the ObjectURLs get cleared.
- Historical/in-flight query performance: You can now use the query list and query details page to analyze historical or in-flight query performance by viewing the memory consumed, the amount of data read, and other information about the query.

## Ranger audit behavior enhancements

Before this release, the Ranger authorization is called for each Impala object; that is, database, table and each column, and this generates a bulky audit for a larger number of columns. This release consolidates the log entries of several columns' accesses into one entry in the same table, which saves space.

### Removing self events

Before this release, some metadata consistency issues lead to query failures because the metadata updates from multiple coordinators could not differentiate between self-generated events and those that are generated by a different coordinator. This issue is resolved now by adding a coordinator flag to each event, and when processing these events we check the coordinator flag to make a decision on whether to ignore the event or not.

### Query hints for table cardinalities

Currently, Impala only uses simple estimation to compute selectivity. For some predicates, the estimation might deviate significantly from the actual value, which leads to a worse query plan. You can now use [a new query hint](#), 'SELECTIVITY', to help specify a selectivity value for a predicate.

### JWT auth for Impala

Authentication is the mechanism to ensure that only specified hosts and users can connect to Impala. To use JWT authentication, you must [configure it in CDP](#) using Cloudera Manager. Clients, such as Impala shell, can then authenticate to Impala using a JWT instead of a username/password.

### Improvements in rolling restart

This release supports the rolling restart of Impala service during the rolling upgrade. However, zero downtime upgrade is not supported yet as Impala is not an HA service, and has singleton components like catalog and statestore. But the [rolling restart](#) has been enhanced to increase the speed by restarting half of the cluster together.

### Using Knox as a proxy

In both CDP public cloud data hubs and private cloud base, clients access Impala through Knox as a proxy for its ability to do SSO. This is the officially encouraged technique, but it requires [setting the parameter in impala-shell](#) -- `http_cookie_names=KNOX_BACKEND-IMPALA` to include the cookie that Knox uses for stateful connections. This config is needed for Active-Active HA to work for Impala.

### Downgrade for Impala

After upgrading to CDP 7.1.9, if you must restore the software back to the pre-upgrade release and preserve the user data to 7.1.8 then do the following for Impala service:

- Stop the running Impala service in CM.
- Rollback the parcel to the older release parcel.
- Start the Impala service.

Note: Suppose time T is the rolling upgrade start time and if you terminate the upgrade by a downgrade then the files created before or after T remain available in HDFS. The files deleted before or after T remain deleted in HDFS.

### Impala Ozone EC support

Impala now supports [reading from Ozone data stored with Erasure Coding](#) (EC). The Ozone EC feature provides data durability and fault tolerance with reduced storage space and ensures data durability similar to the Ratis THREE replication approach. EC can be considered as an alternative to replication.

### Spill to Ozone

You can now use [Ozone as a scratch space](#) for writing intermediate files during large sorts, joins, aggregations, or analytic function operations.

### Ability to create an external table

A user can now [create an external Kudu table](#) pointing to an existing Kudu table if the user is granted the RWSTORAGE privilege on the resource specified by a storage handler URI. Before this release, a user was required

to have the ALL privilege on SERVER to create an external Kudu table for an existing Kudu table. This has been simplified by the introduction of a new type of resource called storage handler URI and a new access type called RWSTORAGE that will be supported by Apache Ranger.

### Ability to create a non-unique primary key for Kudu

Impala now supports [creating a Kudu table with a non-unique primary key](#). When creating a Kudu table, specifying PRIMARY KEY is optional now. If there is no primary key attribute specified, the partition key columns could be promoted as non-unique primary keys if those columns are the beginning columns of the table.

### TPC-DS performance improvements

In this release, the following enhancements are introduced in multiple areas in the planner and executor to improve query performance and to meet the TPC decision support (TPC-DS) benchmark.

- Improve cardinality estimation for [joins involving multiple conjuncts](#).
- Introduced new query options to improve memory [estimation for aggregation nodes](#).
- [Planner changes for CPU usage](#)

This release brings some changes to the query planner to improve parallel sizing and resource estimation. This is done for workload-aware autoscaling and will be available as query options. These additional query options are added for tuning purposes. This new functionality will allow additional customers to enable multi-threaded queries globally for improved performance.

- Impala late materialization of columns

This release [introduces late materialization](#), which optimizes certain queries on Parquet tables by limiting table scanning. Only relevant data is materialized to improve query response.

### Binary support

Impala now supports [BINARY columns](#) for all table formats except Kudu. See the BINARY support topic for more information on using this arbitrary-length byte array data type in CREATE TABLE and SELECT statements.

### ALTER VIEW support

Before this release, altering only the VIEW definition, VIEW name, and owner was supported. Impala now [supports altering](#) the table properties of a VIEW by using the set tblproperties clause.

### BYTES function support

Impala now supports the [BYTES\(\) function](#). This function returns the number of bytes contained in a byte string.

### Resolving ORC columns by names

Before this release, Impala resolved ORC columns by index. In this release, [a query option ORC\\_SCHEMA\\_RESOLUTION](#) is added to support resolving ORC columns by names.

### Retrieving the data file name

Impala now supports including [a virtual column in a standard SELECT statement](#) select INPUT\_\_FILE\_\_NAME from <tablename> to retrieve the name of the data file that stores the actual row in a table.

### Min/Max filtering in Impala

Using Parquet format, you can query to find the [minimum or maximum](#) value for a column within a partition, row group, page, or row.

### Reading and writing Parquet bloom filters

[Bloom filter](#) is a performance optimization feature now available in Impala. This filter tells you, rapidly and memory-efficiently, whether the data you are looking for is present in a file.

### Printing query results in vertical format

Impala-shell now includes a new command option '-E' or '--vertical' to [support printing of query results in vertical format](#).

### Added support for thrift-0.16.0

### Limited support for Hive Generic UDFs

Hive has 2 types of UDFs. This release contains [limited support for the second generation UDFs called GenericUDFs](#). The main limitations are as follows:

- Decimal types are not supported
- Complex types are not supported
- Functions are not extracted from the jar file

GenericUDFs cannot be made permanent. They will need to be recreated every time the server is restarted.

### Reset all query options

UNSET ALL can [unset all](#) query options. This is especially useful when connections are reused, e.g. when a connection pool is used.

## What's New in Apache Kafka

Learn about the new features of Kafka in Cloudera Runtime 7.1.9.

### Rebase on Kafka 3.4.1

Kafka shipped with this version of Cloudera Runtime is based on Apache Kafka 3.4.1. For more information, see the following upstream resources:

Apache Kafka Notable Changes:

- [3.2.0](#)
- [3.3.0 and 3.3.1](#)
- [3.4.0](#)

Apache Kafka Release Notes:

- [3.2.0](#)
- [3.3.0](#)
- [3.3.1](#)
- [3.4.0](#)
- [3.4.1](#)

### Kafka KRaft [Technical Preview]

Apache Kafka Raft (KRaft) is a consensus protocol used for metadata management that was developed as a replacement for Apache ZooKeeper. Using KRaft for managing Kafka metadata instead of ZooKeeper offers various benefits including a simplified architecture and a reduced operational footprint.

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.

For a conceptual overview on KRaft, see [Kafka KRaft](#). For more information on how to set up a cluster with KRaft, see [KRaft setup](#).

### Kafka log directory monitoring improvements

A new Cloudera Manager chart, trigger, and action is added for the Kafka service. These assist you in monitoring the log directory space of the Kafka Brokers, and enable you to prevent Kafka disks from filling up.

The chart is called Log Directory Free Capacity. It shows the capacity of each Kafka Broker log directory.

The trigger is called Broker Log Directory Free Capacity Check. It fires if the capacity of any log directory falls below 10%. The trigger is automatically created for all newly deployed Kafka services, but must be created with the Create Kafka Log Directory Free Capacity Check action for existing services following an upgrade.

The chart and trigger are available on the [Kafka service Status](#) page. The action is available in [Kafka service Actions](#).

### Kafka Connect metrics reporter security configurable in Cloudera Manager

New, dedicated Cloudera Manager properties are introduced for the security configuration of the Kafka Connect metrics reporter. As a result, you are no longer required to use advanced security snippets if you want to secure the metrics reporter and its endpoint. The new properties introduced are as follows:

- Secure Jetty Metrics Port
- Enable Basic Authentication for Metrics Reporter
- Jetty Metrics User Name
- Jetty Metrics Password

A dedicated property to enable TLS/SSL for the metrics reporter is not available. Instead you must select Enable TLS/SSL for Kafka Connect which enables TLS/SSL for the Kafka Connect role including the metrics reporter. For more information regarding these properties, see [Cloudera Manager Configuration Properties Reference](#).

In addition, the Kafka Connect Prometheus Metrics Port property is removed and is replaced by Jetty Metrics Port or Secure Jetty Metrics Port. As a result, the setup steps required to configure Prometheus as the metrics store for SMM are changed. For updated deployment instructions, see [Setting up Prometheus for Streams Messaging Manager](#).

### Single Message Transforms (SMT) plugins for binary conversion

Two Cloudera developed SMT plugins are added. These are the ConvertToBytes and ConvertFromBytes plugins, which you can use to convert binary data to or from the Kafka Connect internal data format.

For more information, see the following resources:

- [Single Message Transforms](#)
- [ConvertFromBytes](#)
- [ConvertToBytes](#)

### Exactly-once semantics (EOS) support for source connectors

EOS support is added for Kafka Connect source connectors. For more information, see [Configuring EOS for source connectors](#).

### Rolling restart checks provide a high cluster health guarantees by default

The default value of the Cluster Health Guarantee During Rolling Restart property is changed from none to healthy partitions stay healthy. This property defines what type of checks are performed during a Rolling Restart on the restarted broker. Each setting guarantees a different level of cluster health during Rolling Restarts. With the none

setting, no checks are performed. This means that in previous versions no guarantees were provided on cluster health by default.

The new default, healthy partitions stay healthy, ensures a high level of guarantees on cluster health. This setting ensures that no partitions go into an under-min-isr state when a broker is stopped. This is achieved by waiting before each broker is stopped so that all other brokers can catch up with all replicas that are in an at-min-isr state. Additionally, the setting ensures that the restarted broker is accepting requests on its service port before restarting the next broker. This setting ignores partitions which are already in an under-min-isr state. For more information, see [Rolling restart checks](#).

### Kafka load balancer is automatically configured with the LDAP handler if LDAP authentication is configured

When a load balancer and LDAP authentication is configured for Kafka, the PLAIN mechanism is automatically added to the enabled authentication mechanisms of the load balancer listener. Additionally, the load balancer is automatically configured to use the `LdapPlainServerCallbackHandler` as the callback handler.

### LDAPS SSL configurations are inherited from the Kafka broker

The SSL configurations of LDAP over SSL (LDAPS) are inherited from the Kafka broker. Previously, the JDK default was used. If the JDK default certificate store contains certificates which were used to setup SSL connection to LDAP, it should be imported to the broker stores.

### Aliases for Kafka CLI tools

Aliases are added for the `kafka-storage.sh`, `kafka-cluster.sh`, and `kafka-features.sh` command line tools. These tools can now be called globally with `kafka-storage`, `kafka-cluster`, and `kafka-features`.



**Important:** Not all tools are fully supported and their use is limited. For more information, see [Unsupported command line tools](#).

### Multi-level rack awareness

The rack awareness capabilities of Kafka are improved to support multi-level cluster topologies. As a result, brokers can now be configured to run in a multi-level rack-aware mode. If this mode is enabled, the brokers provide multi-level rack awareness guarantees. These guarantees ensure that topic partition replicas are spread evenly across all levels of the physical infrastructure. For example, in a two-level hierarchy with Data Centers on the top level and racks on the second level, brokers will evenly spread replicas among both available DCs and racks.

The new mode is compatible with follower fetching. If multi-level mode is enabled, a compatible replica selector class is automatically installed. This implementation enables consumers (if configured), to fetch Kafka messages from the replica that is closest to them in the multi-level hierarchy.

Additionally, when Cruise Control is deployed on the cluster, the standard rack-aware goals in Cruise Control's configuration are replaced with a multi-level rack-aware goal. This goal ensures that Cruise Control optimizations do not violate the multi-level rack awareness guarantees. This goal is currently downstream only, available exclusively in Cloudera distributed Cruise Control. For more information, see the following resources:

- [Kafka rack awareness](#)
- [Configure Kafka rack awareness](#)
- [Setting capacity estimations and goals](#)

### AvroConverter support for Kafka Connect logical types

The `AvroConverter` now converts between Connect and Avro temporal and decimal types.



### Connector configurations must by default override the `sasl.jaas.config` property of the Kafka clients used by the connector

The Require Connectors To Override Kafka Client JAAS Configuration Kafka Connect property is now selected by default. This means that connector configurations must by default contain a `sasl.jaas.config` entry with an appropriate JAAS configuration that can be used to establish a connection with the Kafka service.

### Connect JAAS enforcement now applies to non-override type Kafka client configs

When the Require Connectors To Override Kafka Client JAAS Configuration property is selected, the `consumer.sasl.` and `producer.sasl.` configurations are not emitted into the Connect worker configurations anymore. Additionally, the keytab name is randomized and the `${cm-agent:keytab}` references in the Connector configurations will stop working.

### Kafka Connect now supports Kerberos auth-to-local (ATL) rules with SPNEGO authentication

Kafka Connect now uses the cluster-wide Kerberos auth-to-local (ATL) rules by default. A new configuration property called Kafka Connect SPNEGO Auth To Local Rules is introduced. This property is used to manually specify the ATL rules. During an upgrade, the property is set to `DEFAULT` to ensure backward compatibility. Following an upgrade, if you want to use the cluster-wide rules, clear the existing value from the Kafka Connect SPNEGO Auth To Local Rules property.

### Debezium connector updates

The following updates related to Debezium connectors are introduced:

- All Debezium connectors shipped with Cloudera Runtime are upgraded to version 1.9.7.  
Existing instances of the connectors are automatically upgraded to the new version during cluster upgrade. Deploying the previously shipped version of the connector is not possible.
- The Debezium Db2 Source connector is introduced and is available for deployment.

For more information see [Kafka Connectors in Runtime](#) or the [Debezium documentation](#).

### Parquet support for the S3 Sink connector

Version 2.0.0 of the S3 Sink connector is released. The connector now supports Parquet as an output file data format. The following property changes are made to support Parquet:

- A new property, Parquet Compression Type, is added.  
This property specifies the compression type used for writing Parquet files. Accepted values are `UNCOMPRESSED`, `SNAPPY`, `GZIP`, `LZO`, `BROTLI`, `LZ4`, and `ZSTD`.
- The Output File Data Format property now accepts Parquet as a value.

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 [S3 Sink connector](#) and [S3 Sink properties reference](#).

### Support schema ID encoding in the payload or message header in Stateless NiFi connectors

The Kafka Connect connectors powered by Stateless NiFi that support record processing are updated to support content-encoded schema references for Avro messages. These connectors now properly support integration with Schema Registry and SMM.

This improvement introduces the following changes in the affected connectors.

**A new value, HWX Content-Encoded Schema Reference, is introduced for the Schema Access Strategy property**

If this value is set, the schema is read from Schema Registry, and the connector expects that the Avro messages contain a content-encoded schema reference. That is, the message contains a schema

reference that is encoded in the message content. The new value is introduced for the following connectors:

- ADLS Sink
- HDFS Sink
- HTTP Sink
- Influx DB Sink
- JDBC Sink
- JDBC Source
- Kudu Sink
- S3 Sink

**The Schema Write Strategy property is removed from the following connectors**

- ADLS Sink
- HDFS Sink
- S3 Sink
- InfluxDB Sink

**A new property, Avro Schema Write Strategy is introduced**

This property specifies whether and how the record schema is attached to the output data file when the format of the output is Avro. The property supports the following values:

- Do Not Write Schema: neither the schema nor reference to the schema is attached to the output Avro messages.
- Embed Avro Schema: the schema is embedded in every output Avro message.
- HWX Content-Encoded Schema Reference: a reference to the schema (identified by Schema Name) within Schema Registry is encoded in the content of the outgoing Avro messages.

This property is introduced for the following connectors:

- ADLS Sink
- HDFS Sink
- S3 Sink
- SFTP Source
- Syslog TCP Source
- Syslog UDP Source



**Note:** With the exception of InfluxDB Sink, this property replaces Schema Write Strategy in connectors where Schema Write Strategy was previously available.

**The minor or major version of all affected connectors is updated**

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 the documentation for each connector in [Kafka Connectors in Runtime](#) and [Streams Messaging Reference](#).

## What's New in Kerberos

Learn about the new features of Kerberos in Cloudera Runtime 7.1.9

### 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 Key Trustee Server

Learn about the new features of Key Trustee Server in Cloudera Runtime 7.1.9

- Postgres 14.2 is the embedded database version for Key Trustee Server in 7.1.9.
- Key Trustee Server now supports RHEL 8.8, RHEL 8.8 FIPS and RHEL 8.6. RHEL 9 is not supported for KTS.

## What's New in Apache Knox

Learn about the new features of Knox in Cloudera Runtime 7.1.9.

### Token-based authentication in Private Cloud

You can now use the Knox homepage to generate and manage Knox Gateway tokens for Cloudera Data Platform.

For more information, see [Knox Gateway token integration](#).

### Knox load balancing for Oozie and Impala

Oozie and Impala now support Knox load balancing.

For more information, see [Load balancing for Apache Knox](#).

### Ozone-Knox Integration

Ozone and Knox are now integrated to work together.

For more information, see [Ozone Knox integration](#).

### Knox Zero Downtime Upgrades and Rolling Restarts

Knox now supports rolling upgrades and rolling restarts.

For more information, see [Zero Downtime upgrade](#).

## What's New in Apache Kudu

Learn about the new features of Kudu in Cloudera Runtime 7.1.9.

### Kudu JWT support and proxy support

JWT authentication is an alternative to Kerberos authentication, and you can use it in situations where Kerberos authentication is not a viable option but authentication is required nevertheless. For more details, see [JWT authentication for Kudu](#).

It is now possible to separate the internal and the external traffic in a Kudu cluster while providing the connectivity for Kudu clients running in external networks where the internal traffic is never routed through a proxy's or a loadbalancer's endpoint. Essentially, it allows for the internal traffic (for example, the traffic between tablet servers and masters) to bypass advertised RPC addresses, using alternative addresses for inter-cluster communications. For more details, see [Proxied RPCs in Kudu](#).

### Auto-incrementing column

Introduced auto-incrementing column. These columns are populated on the server side with a monotonically increasing counter. The counter is local to every tablet; for example, each tablet has a separate auto incrementing counter.

Kudu now supports experimental non-unique primary key. When a table with non-unique primary key is created, an auto-incrementing column named `auto_incrementing_id` will be added automatically to the table as the key column. The non-unique key columns and the auto-incrementing column together form the effective primary key (see, KUDU-1945). For more details, see [Non-unique primary key index](#).

### Auto-leader rebalancing

An experimental feature is added to Kudu that allows it to automatically rebalance tablet leader replicas among tablet servers. The background task can be enabled by setting the `--auto_leader_rebalancing_enabled` flag on the Kudu masters (see, KUDU-3390).

### Immutable column

Introduced immutable column. It is useful to define such a column which represents a semantically constant entity (see, KUDU-3353).

### Added sanity check to detect wall clock jumps

Added a sanity check to detect strange jumps in wall clock readings. The idea is to rely on the readings from the `CLOCK_MONOTONIC_RAW` clock captured along with the wall clock readings. A jump should manifest itself in a big difference between the wall clock delta and the corresponding `CLOCK_MONOTONIC_RAW` delta. If such a condition is detected, then `HybridClock::NowWithErrorUnlocked()` dumps diagnostic information about clock NTP synchronisation status and returns `Status::ServiceUnavailable()` with appropriate error message.

As a part of this changelist, the following new flags are introduced:

- `--wall_clock_jump_detection`

This is to control the newly introduced sanity check for readings of the wall clock. Acceptable values are `auto`, `enabled`, and `disabled`. It is set to `auto` by default, which means that the sanity check for timestamps is enabled if the process detects that it is running on a VM in Azure cloud.

- `--wall_clock_jump_threshold_sec`

This is to control the threshold (in seconds) for the difference in deltas of the wall clock's and `CLOCK_MONOTONIC_RAW` clock's readings. It is set to 900 (15 minutes) by default.

### Kudu multi-master config change

You can now remove or decommission the unwanted master role instances through Cloudera Manager. Also, you can recommission any decommissioned master role instance in a multi-master deployment. For more information, see [Remove Kudu masters through Cloudera Manager](#).

### Kudu Range-aware Data Placement

Kudu places new tablet replicas using an algorithm which is both range and table aware. This algorithm helps to avoid hotspotting that occurs if many replicas from the same range are placed on the same few tablet servers. Hotspotting causes tablet servers to be overwhelmed with write or read requests and can result in increased latency for these requests. To avoid hotspotting, this algorithm avoids targeting the same set of tablet servers for a set of replicas created in parallel. Rather, it spreads the replicas across multiple tablet servers. For more information, see [Range-aware replica placement in Kudu](#).

### Kudu replication factor

A new tool, `kudu table set_replication_factor`, is added to dynamically modify the replication factor of an existing table. You can now adjust the replication factor for existing tables without needing to recreate them, providing more flexibility in managing how data is replicated across your cluster. The tool immediately updates the table metadata in the master, and the master asynchronously applies the new replication factor. You can monitor the progress by running the `ksck` command.

## What's New in Livy

Learn about the new feature of Livy in Cloudera Runtime 7.1.9.

### High Availability support added for Livy

Livy now supports high availability. If there are more than one Livy Server in the cluster, high availability is automatically enabled.

For more information, see [Livy high availability support](#).

## What's New in Navigator Encrypt

Learn about the new features of Navigator Encrypt in Cloudera Runtime 7.1.9

- NavEncrypt now supports Linux Distributions RHEL 9.1, RHEL 8.8, RHEL 8.8 FIPS, and SLES15-SP4.
- Integration of NavEncrypt with Ranger KMS - When configuring NavEncrypt to interoperate with Ranger KMS, be advised that this will require NavEncrypt to use kerberos. For more information, refer to topics at : [Navigator Encrypt Overview](#).
- When the NavEncrypt kernel module emits a system log message, the originator's name is now navencryptfs instead of navencrypt.

## What's New in Apache Oozie

Learn about the new features of Apache Oozie in Cloudera Runtime 7.1.9.

### Spark 3 support in Oozie

Oozie introduced the new Spark 3 based Spark 3 actions. For more information, see [Spark 3 support in Oozie](#).

### Make hive-site.xml, hbase-site.xml and sqoop-site.xml available for all Oozie actions

Now Oozie automatically copies the hive-site.xml, hbase-site.xml, and sqoop-site.xml to all action's classpath. For more information, see [Oozie and client configurations](#).

### Improve Oozie's app state action checking

Enhanced Oozie's action state checking, to immediately query for running applications right after start-up.

### Oozie should upload and use the config files from sqoop-conf/managers.d when available

Previously, Oozie did not honor Sqoop's managers.d configurations and extra connector Jars from the lib folder, but now both are automatically available in Oozie's Sqoop action, allowing users to seamlessly utilize connectors like the Sqoop Teradata connector without the need for manual configuration updates or copying Jars to the Workflow's lib folder.

### Oozie should not rely on its LoadBalancer internally

Oozie will no longer use the LoadBalancer to issue a callback notification, but instead it will try all available Oozie instances one-by-one. If the callback succeeded against one of the Oozie instances, then we will not try the other ones. This way the LoadBalancer will not be used for such purposes.

Cloudera Manager will provide the address of all Oozie server instances as a configuration to all Oozie instances. This will be then used by Oozie's callback mechanism so that instead of making the callback through the LoadBalancer in HA mode, the callback will be attempted through each Oozie instance, and if one of them succeeds, then we stop. This way we'll no longer use the LoadBalancer, and make the callback mechanism safer by not having a middle-man.

### Handle Sqoop Teradata Connector parcels installation and configuration for Oozie

When you install a Sqoop Teradata connector parcel, Cloudera Manager will automatically make the necessary Jars and configuration available to Oozie's Sqoop action.

### TCPS support for Oozie with Oracle DB backend

When it comes to configuring database connections, simply providing a hostname, port, username, and password may not be sufficient. In order to optimize Oozie's database connection, you might need to manually construct lengthy connection and configuration strings using safety-valve settings. To simplify this process and enable finer control over Oozie's database connection, you can use several enhancements. For more information, see [Fine-tuning Oozie's database connection](#).

### JDK 17 support

In supporting Java 17, certain applications which are executed with Oozie might require reflective access to internal Java classes, packages, or modules. To enable reflective access, you need to use the add-opens Java parameter. For more information, see [Oozie Java-based actions with Java 17](#).

## What's New in Apache Ozone

Learn about the new features of Apache Ozone in Cloudera Runtime 7.1.9.

### Snapshot support in Ozone

Learn about different scenarios where you can use snapshots, the snapshot APIs that are available for use, and the snapshot architecture.

For more information, see [Working with snapshots in Ozone](#).

### Erasure Coding Enhancements

The Ozone Erasure Coding feature provides data durability and fault-tolerance along with reduced storage space and ensures data durability similar to Ratis THREE replication approach.

For more information, see [Erasure Coding](#).

### SCM Certificate Rotation

Configuring security and issuing certificates to Storage Container Managers (SCMs) along with Ozone Managers (OMs) and the DataNodes ensures secured communication in the Ozone cluster.

For more information, see [Configuring security for Storage Container Managers in High Availability](#).

### Master node decommissioning in Ozone

This feature helps you to decommission Ozone Manager (OM) and Storage Container Manager (SCM).

For more information, see [Master node decommissioning in Ozone](#).

### Recon Heat Map

You can access the heatmap feature as an administrator to read or view the most accessed volumes, buckets, and top 100 keys across Apache Ozone.

For more information, see [Recon Heat Map](#).

### Managing Ozone quota

The Ozone shell is the primary command line interface for managing the quota of volumes and buckets.

For more information, see [Managing Ozone quota](#).

### Ozone Knox Integration

Ozone and Knox are now integrated to work together.

For more information, see [Ozone Knox integration](#).

### Configuration options for Oozie to work with Ozone storage

Oozie supports Ozone storage along with HDFS.

Apache Ozone is a highly scalable next-gen object store available on the CDP Private Cloud Base cluster which enables you to optimize storage for big data workloads.

For more information, see [Configuration options for Oozie to work with Ozone storage](#).

### Ozone HttpFS support

Ozone now supports HttpFS Gateway. This allows Ozone to integrate with other tools through REST APIs.

For more information, see [Ozone HttpFS support](#).

### Configuration options for Impala to work with Ozone File System

Learn how Ozone can work with Impala.

You can use Impala to query data files that reside on Apache Ozone distributed storage, rather than in HDFS.

For more information, see [Configuration options for Impala to work with Ozone File System](#).

### Determining the threshold

Container Balancer balances the utilization of DataNodes in a cluster using the Threshold. You can now determine the threshold value before configuring the required parameters.

For more information, see [Determining the threshold](#).

### Volume and bucket management using ofs

This is a new improvement to the existing feature where you can now use \_ in naming volume and bucket.

For more information, see [Volume and bucket management using ofs](#).

### Ozone volume scanner

The Ozone Volume Scanner feature enables to detect any disk failures on the DataNodes. Learn how you can configure the frequency of volume scans that can detect disk failures and how to handle volume failures.

For more information, see [Ozone volume scanner](#).

### Ozone OMDBInsights

The Ozone Manager Database Insights feature helps you view the container mismatch information, open keys, keys pending for deletion, and deleted container keys.

For more information, see [Ozone OMDBInsights](#).

## What's New in Apache Phoenix

Learn about the new features of Apache Phoenix in Cloudera Runtime 7.1.9.

### 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

The following new features and enhancements are generally available for Ranger customers in Cloudera Runtime 7.1.9:

### Ranger Replication

You can create Ranger replication policies in CDP Private Cloud Base Replication Manager. The Ranger replication policies migrate the Ranger policies, roles, and tags for HDFS, Hive, and HBase services between Kerberos-enabled CDP Private Cloud Base 7.1.9 or higher clusters using Cloudera Manager 7.11.3. It can also migrate Ranger audit logs in HDFS. For more information, see the topics at: [Ranger replication policies](#).

### Ranger Usersync option to update group memberships when same users and groups are synced from multiple sync sources

Ranger Usersync now provides an option for customers to treat users/groups from multiple sync sources as the same for updating group memberships. For more information, see the updated topic: [Configuring Usersync to sync directly with LDAP/AD](#).

### HA support for Ranger Tag Sync/User Sync

Ranger now supports high availability for Ranger Tag Sync/User Sync. Configuring high availability adds another instance of each role to an additional host, which host continues to run the features if the default host fails. For more information, see [Configuring Ranger Usersync and Tagsync High Availability](#).

### New Ranger API to collect metrics in Ranger Admin

Ranger now provides two APIs to fetch ranger admin metrics. One returns a response in JSON format and the other returns a response in prometheus-compatible format. For more information, see [Ranger Admin Metrics API](#).

### New Ranger APIs to import/export roles in Ranger Admin

Ranger now includes APIs to import and export roles. For more information, see [Ranger REST API documentation](#).

### Ranger HDFS plugin option to view permissions through getfacl interface when Ranger RMS (Hive-HDFS ACL Sync) is enabled

You can configure the Ranger HDFS plugin to view user access permissions in a manner similar to the HDFS `getfacl` command after migrating from CDH to CDP. This change is just a way to see the permissions. There is NO change in the way Ranger RMS (Hive-HDFS ACL Sync) enforces permissions. For more information, see [Configuring HDFS plugin to view permissions through getfacl interface](#).

### Ranger RMS support for Ozone

In CDP 7.1.9, Ranger RMS will support authorization for Ozone storage locations. RMS for Ozone will co-exist with Hive-HDFS ACL sync and provide authorization for both HDFS and Ozone file systems. For more information, see the updated topics and examples throughout: [About Ranger RMS for Ozone](#).



### Add support for enabling audit file accumulation

You can enable and configure alerts for Ranger plugin-supported services through Cloudera Manager. Such alerts notify when audit spool files accumulate in the spool directories for Solr and HDFS. For more information, see [Configuring audit spool alert notifications](#).

### Add support for additional methods in RangerKafkaAuthorizer

RangerKafkaAuthorizer includes ACL APIs that refer to Ranger Policies when these commands are executed. Ranger relies on the grant, revoke and policy engine APIs to cater the needed functionality. For more information, see [Kafka ACL APIs support](#).

### Add APIs to support force deletes of external users and groups from Ranger db

A Ranger database may (over)-populate with user and group records. To aid in removal of unnecessary users/groups, customers may use this feature to delete specific external user/groups or even all external users/groups if required. For more information, see [Force deletion of external users and groups from the Ranger database](#).

### Performance and Function Improvements

- Provide workaround for Ranger RMS customers who may experience intermittent high RPC queue and processing time. For more information, see [Ranger RMS field issues - HDFS latency](#).

## Ranger KMS

Learn about the new features of Ranger KMS in Cloudera Runtime 7.1.9

### Ranger KMS Key Migration

Ability to migrate keys from Key Trustee Server to Ranger KMS DB. For more information, refer to : [Migrating keys from Key Trustee Server to Ranger KMS](#).

### Pass JVM options to Ranger KMS KTS services

Added field within RANGER KMS to add items to the JAVA\_OPTS environment variable to enable better debugging and tuning. For more information , refer to: [How to pass JVM options to Ranger KMS KTS services](#).

### Ranger KMS Health Metrics

Additional Ranger KMS Server health metrics have been added to Cloudera Manager. For more information, refer to [Ranger KMS Server Metrics](#).

### Integration of Ranger KMS with Luna 10.5 HSM

How to integrate Cloudera Ranger Key Management System (KMS) software with the Luna 10.5 HSM appliance supplied by SafeNet. For more information, refer to : [Set up Luna 10.5 HSM Client for Ranger KMS](#).

### Ranger KMS Ozone support

Ranger KMS Ozone support is available in 7.1.9. For more information, refer to [Configuring Transparent Data Encryption for Ozone](#).

### Linux distribution support

Ranger KMS now supports Linux Distributions RHEL 9.1, RHEL 8.8, and RHEL 8.8 FIPS

### Ranger KMS supports connections to databases secured using TLS 1.2 and TCPS

Ranger KMS can connect to TLS-enabled MySQL, MariaDB, or PostgreSQL databases and TCPS-enabled Oracle database. To connect to a TLS/TCPS-enabled database while adding the Schema Registry service to a cluster, see

[Configure TLS 1.2 for Ranger KMS](#). You can also enable TLS/TCPS on an existing database and then configure Schema Registry to connect to it. See [Set up and configure TLS 1.2 for Ranger KMS](#). For more information about Oracle TCPS, see [How to connect CDP components to a TCPS-enabled Oracle database](#).

## What's New in Schema Registry

Learn about the new features for Schema Registry in Cloudera Runtime 7.1.9.

### Schema Registry supports connections to databases secured using TLS 1.2 and TCPS

Schema Registry can connect to TLS-enabled MySQL, MariaDB, or PostgreSQL databases and TCPS-enabled Oracle database. To connect to a TLS/TCPS-enabled database while adding the Schema Registry service to a cluster, see [Configure TLS 1.2 for Schema Registry](#). You can also enable TLS/TCPS on an existing database and then configure Schema Registry to connect to it. See [Set up and configure TLS 1.2 for Schema Registry](#). For more information about Oracle TCPS, see [How to connect CDP components to a TCPS-enabled Oracle database](#).

### Schema Registry instances behind load balancer

You can now use load balancer in front of Schema Registry instances. It is very common to have multiple instances of the same application and have a load balancer in front of them. This can be useful for failover reasons in HA environments, and it can also help sharing the load between instances. You can also use load balancer in front of Schema Registry instances in an environment with Kerberos or SSL enabled.

### AvroConverter support for KConnect logical types

AvroConverter now converts between Connect and Avro temporal and decimal types.

### Support for alternative jersey connectors in SchemaRegistryClient

`connector.provider.class` can be configured in Schema Registry Client. If it is configured, `schema.registry.client.retry.policy` should also be configured to be different than default.

This also fixes the issue with some third party load balancers where the client is expected to follow redirects and authenticate while doing that.

### Schema Registry with Knox uses round-robin load balancing

When multiple instances of Schema Registry are running, Knox uses round-robin to forward the requests.

### Upgraded Avro version to 1.11.1

Avro got upgraded from version 1.9.1 to 1.11.1.

### KafkaAvroSerializer and KafkaAvroDeserializer improvements

#### KafkaAvroSerializer and KafkaAvroDeserializer can now handle null values without Avro

The KafkaAvroSerializer and KafkaAvroDeserializer now support a configuration property called `null.passthrough.enabled`, which is false by default. If enabled, null data is handled as null, and no schema is sent to Schema Registry. This behavior enables client applications to write tombstone messages into compact topics. The KafkaAvroDeserializer also handles null values by returning null without any regards to the schema.

#### Support deserialization when the topic and schema names don't match

From now on, the KafkaAvroDeserializer uses the schema version's ID in the Avro byte stream to access the actual schema and disregards schema names.

#### Logical types conversion for the KafkaAvroSerliazier and KafkaAvroDesrializer

The KafkaAvroSerializer and KafkaAvroDeserializer can now properly handle and convert Avro logical types at a record level. This means that if you have a record that has a field with a built-in Avro logical type (for example a BigDecimal field with BYTES type and decimal logical type), you

can now properly serialize the records. After deserialization, a `GenericRecord` is returned, including the typed `BigDecimal` field, instead of a `ByteBuffer`. Logical type conversion can be enabled using the `logical.type.conversion.enabled` property. This property is set to `false` by default for backward compatibility.

For more information, see the following resources:

- [KafkaAvroDeserializer properties reference](#)
- [KafkaAvroSerializer properties reference](#)

### Principal mapping rules can be defined without quotes

The `SSL Client Authentication Mapping Rules` (`schema.registry.ssl.principal.mapping.rules`) property now accepts rules that are defined without quotes. As a result, when adding multiple rules, you no longer need to enclose each rule in quotes.

### Modules section are removed from the `registry.yaml` configuration structure

In previous versions, the `registry.yaml` configuration file contained a `modules` section. This section was used to list pluggable modules that extended Schema Registry's functionality. However, modules were never fully supported and have been removed in a previous release. The `modules` section in `registry.yaml` was kept for backwards compatibility. Starting with this version, the `modules` section is removed by default from `registry.yaml`.

## What's New in Apache Solr

Learn about the new features of Solr in Cloudera Runtime 7.1.9.

- 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](#).
- Solr now supports rolling upgrades. This means that rolling upgrades are available from release 7.1.9 to higher. Upgrades to release 7.1.9 still involve service downtime.
- 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 Spark in Cloudera Runtime 7.1.9.

### CDS 3.3 powered by Apache Spark

The default Spark runtime in Cloudera Runtime version 7.1.9 for CDP Private Cloud Base is Spark 2.4.8 which is deprecated. CDS 3.3 is an add-on parcel that can be installed to provide support for Spark 3.3. Additionally, CDS

3.3 is certified for Cloudera Runtime version 7.1.9 for CDP Private Cloud Base based on Spark version 3.3.2 and contains all the feature content of that release.

- Support for virtual clusters powered by Apache Spark 3 is now available.
- Spark 2 is deprecated in Cloudera Runtime 7.1.9, therefore 7.1.9 is the last Cloudera Runtime release where Spark 2 is supported.
- Support for Hive Warehouse Connector (HWC) - that is, Hive managed ACID tables (Direct Reader and JDBC mode)
- The following functionalities are not currently supported:
  - Deep analysis (visual profiler)
  - Phoenix Connector
  - SparkR

See [Running Apache Spark 3 applications](#) and [Deprecation Notices for Spark 2](#).

### Spark 3 support in Oozie

Oozie introduced the new Spark 3 based Spark 3 actions. For more information, see [Spark 3 support in Oozie](#)

### Spark History Server with High Availability

You can configure the load balancer for Spark History Server (SHS) to ensure high availability, so that users can access and use the Spark History Server UI without any disruption. Learn how you can configure the load balancer for SHS and the limitations associated with it. For more information, see [Using Spark History Servers with high availability](#).

## What's New in Sqoop

Learn about the new features of Sqoop in Cloudera Runtime 7.1.9.

To access the latest Sqoop documentation on Cloudera's documentation web site, go to [Sqoop Documentation 1.4.7.7.1.6.0](#).

### Sqoop enhancements to the Hive import process

This release introduces several Sqoop enhancements that enable you to configure how Sqoop imports data from relational databases into Hive. With these enhancements, you can now specify custom Beeline arguments, define custom Hive JDBC arguments, choose how tables are created in Hive using custom CREATE TABLE statements, and configure custom Hive table properties. The changes allow users to control the imported data according to their specific requirements. For more information, see [Sqoop enhancements to the Hive import process](#).

### Sqoop Teradata Connector support for ORC file format

A new version of Cloudera Connector Powered by Teradata version 1.8.5.1c7 is released which includes ORC support in the Sqoop-Connector-Teradata component. You can use Teradata Manager to import data from the Teradata server to Hive in ORC format. For more information, see [Cloudera Connector Powered by Teradata Release Notes](#)

### Discontinued maintenance of direct mode

The Sqoop direct mode feature is no longer maintained. This feature was primarily designed to import data from an abandoned database, which is no longer updated. 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.

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`.

## What's New in Streams Replication Manager

Learn about the new features of Streams Replication Manager in Cloudera Runtime 7.1.9.

### Prefixless replication with the `IdentityReplicationPolicy`

Full support, including replication monitoring, is introduced for the `IdentityReplicationPolicy`. Unlike the `DefaultReplicationPolicy`, this policy does not rename remote (replicated) topics on target clusters. That is, the topics that you replicate will have the same name in both source and target clusters. This replication policy is recommended for deployments where SRM is used to aggregate data from multiple streaming pipelines. Alternatively, this replication policy can also be used if the deployment requires MirrorMaker1 (MM1) compatible replication.

Prefixless replication is enabled in Cloudera Manager with the `Enable Prefixless Replication` property. This property configures SRM to use the `IdentityReplicationPolicy` and enables internal topic based remote topic discovery, which is required for replication monitoring.

#### Limitations:

- Replication loop detection is not supported. As a result, you must ensure that topics are **not** replicated in a loop between your source and target clusters.
- The `/v2/topic-metrics/{target}/{downstreamTopic}/{metric}` endpoint of SRM Service v2 API does not work properly with prefixless replication. Use the `/v2/topic-metrics/{source}/{target}/{upstreamTopic}/{metric}` endpoint instead.
- The replication metric graphs shown on the **Topic Details** page of the SMM UI do not work with prefixless replication.



**Important:** If you have been using the `IdentityReplicationPolicy` in a previous version of Cloudera Runtime, ensure that you transition your configuration and set the `IdentityReplicationPolicy` with `Enable Prefixless Replication`. If you do not transition your configuration, replication monitoring will not function.

For more information, see

- [Streams Replication Manager replication flows and replication policies](#)
- [Enabling prefixless replication](#)
- [Step 9: Complete Post-Upgrade steps for upgrades to CDP Private Cloud Base](#)

### Internal topic based remote topic discovery

From now on, SRM uses an internal Kafka topic to keep track of remote (replicated) topics. Previously, SRM relied on the naming conventions (prefixes) used by the `DefaultReplicationPolicy` to discover and track remote topics.

This feature enables SRM to provide better monitoring insights on replications. Additionally, if the feature is enabled, SRM is capable of providing replication monitoring even if a replication policy different than the `DefaultReplicationPolicy` is in use. Most notably, this enables replication monitoring when SRM is configured for prefixless replication with the `IdentityReplicationPolicy`.

This feature is enabled in Cloudera Manager by selecting the `Remote Topics Discovery With Internal Topic` property. The property is selected by default on newly deployed clusters, but must be enabled manually for existing clusters after an upgrade. Cloudera recommends that you enable this feature no matter what replication policy you are using.

For more information, see [Streams Replication Manager remote topic discovery](#)

### Improvements related to raw metric collection and aggregation

#### New metric replication-records-lag

The SRM Service now reports a new metric on its REST API called replication-records-lag. This metric provides information regarding the replication lag based on offsets. The metric is available both on the cluster and the topic level.

### Raw metrics are compressed using LZ4

SRM internal metric producers from now on use LZ4 compression by default. LZ4 was chosen as it provides the best combination in terms of compression speed and performance. As a result, Cloudera recommends that you use LZ4. If required, however, you can change the compression by doing the following:

1. Add the following configuration entries to Streams Replication Manager's Replication Configs:

```
workers.cloudera.metrics.reporter.producer.compression.type=[***COMPRESSION***]

connectors.cloudera.metrics.reporter.producer.compression.type=[***COMPRESSION***]
```

2. Add the following to Additional Configs For Streams Application Running Inside SRM Service:

```
producer.compression.type=[***COMPRESSION***]
```

For more information regarding, metrics, monitoring, as well as raw metric collection and aggregation, see [Streams Replication Manager monitoring and metrics](#).

### Improved SRM logging

The logging capabilities of SRM are improved. The following improvements have been introduced:

- Kafka clients created by the internal connectors of SRM reference the replication flow the connectors are a part of ([KAFKA-14838](#) backport).
- SRM now includes references to the replication flow in the log context of its internal connectors.

These changes enable differentiation between the logs associated with each replication flow.

### Metrics and health checks for the status processor Streams application in SRM Service

SRM Service health tests now show the state of the Connect status processor Streams application.

### SRM topic creation timeout increased

Streams Replication Manager internal topic creation timeout property defaults are increased to 20s to tolerate intermittent issues at startup.

## What's new in Streams Messaging Manager

Learn about the new features for Streams Messaging Manager in Cloudera Runtime 7.1.9.

### UI updates

Various improvements are introduced for the SMM UI. The notable changes are as follows:

#### Data Explorer

- When you view Avro data in the Data Explorer, logicalTypes are converted by default. That is, instead of showing the underlying type, (for example, byte) the Data Explorer displays proper deserialized values.
- Avro messages are now pretty printed when you open them using the Show More option.
- The modal window that you use to view messages now includes a copy to clipboard button if the message you are viewing is long.

- A  (Refresh) option is added next to the FROM OFFSET field. This option refreshes the partition offset range and fetches the latest messages.
- You can now view JSON output in pretty printed format by selecting the JSON Pretty Print deserializer.
- A new drop-down, ISOLATION LEVEL:, is added to the Data Explorer.

The drop-down configures the isolation.level property of the Kafka consumer that the Data Explorer uses. The isolation level is also configurable when using the REST API with the consumerIsolationLevel parameter of the /api/v1/admin/topics/{topicName}/partition/{partitionId}/payloads endpoint. The accepted values are read\_committed and read\_uncommitted. The default value is read\_uncommitted.

### Kafka Connect

- Hovering over the status icons of connector tasks now displays the status text instead of the name of the icon.
- The Add missing configurations option now populates missing properties with default values.
- Adding flow.snapshot into a key field of a password type property clears password placeholders.
- An error page is displayed if you navigate to a connector that does not exist.
- A new option, Add, is added to the **Import a Connector config...** modal. This option enables you to import connector configuration properties without overriding existing properties.
- Search and autocomplete are now available for connector property keys. In addition, you can now filter property keys based on their group and importance.
- An error message is added that notifies you if validation errors are found for properties that are currently filtered.
- A Reset Filters option is added. This option resets all search filters.
- Three new actions are added that modify the configuration as a whole. The options are Remove all, Reset, and Export. These actions are available in a new Actions drop-down.
- The Import Connector Configuration... option is moved to the Actions drop-down and is renamed to Import.
- The **Deployment Status** modal now correctly displays the status of the deployment process.
- If available, the display names of configuration property keys are displayed above the property key.
- The **Connector Setup** wizard is updated. Connector (template) selection and connector configuration now happen on separate pages of the wizard.
- A  (Help) option is added that provides detailed information about each property key. The  icon is only available for properties that have their metadata (description, type, group, and so on) defined.
- Types can be specified for properties.

For more information regarding the various new features and options related to Kafka Connect, see [Managing and monitoring Kafka Connect using Streams Messaging Manager](#).

### Other

- The style of SMM UI is updated. This update includes various changes to the colors, fonts, and overall style of the UI.
- You can now increase the number of partitions of a topic. The option is available on the **CONFIGS** tab on the **Topic Details** page. Decreasing the partitions of a topic is not possible.
- The SMM UI contains **Brokers** and **Topics** pages where records contain broker or topic specific partition lists and their profile pages as well. All partition list columns are now sortable.
- Log-size related information is now displayed about brokers, topics, and partitions. Furthermore, warning messages appear when log directory related errors are reported by Kafka.
- The **Cluster Replications** tab now also shows the replication-records-lag metric.



### Changes in Prometheus setup and configuration

In this version of Cloudera Runtime, a number of improvements are introduced related to how Kafka Connect exposes its metrics. Most notably:

- Kafka Connect is now capable of securing its metric reporter with TLS/SSL and Basic Authentication
- The default port of the Kafka Connect metric reporter is changed from 28084 to 28806 (unsecure) or 28807 (secure).
- The Kafka Connect Prometheus Metrics Port property is removed and is replaced by Secure Jetty Metrics Port and Jetty Metrics Port.

As a result, the setup steps required to configure Prometheus as the metrics store for SMM are changed. For updated deployment instructions, see [Setting up Prometheus for Streams Messaging Manager](#).

If you already use Prometheus with SMM, you must make changes to your Prometheus configuration following a cluster upgrade and update the Kafka Connect port that Prometheus connects to. If configuration is not done, Kafka Connect metrics will no longer be available in SMM. For exact steps, see [Step 9: Complete Post-Upgrade steps for upgrades to CDP Private Cloud Base](#).

### SMM supports connections to databases secured using TLS 1.2 and TCPS

SMM can connect to TLS-enabled MySQL, MariaDB, or PostgreSQL databases and TCPS-enabled Oracle database. To connect to a TLS/TCPS-enabled database while adding the SMM service to a cluster, see [Configure TLS 1.2 for Streams Messaging Manager](#). You can also enable TLS/TCPS on an existing database and then configure SMM to connect to it. See [Set up and configure TLS 1.2 for Streams Messaging Manager](#). For more information about Oracle TCPS, see [How to connect CDP components to a TCPS-enabled Oracle database](#).

### SMM internal Kafka topics are created with a replication factor of 3

From now on the \_\_smm\* internal SMM topics are created with a replication factor of 3. This change is only true for newly deployed clusters. The replication factor is not updated during the upgrade. Cloudera recommends that you increase the replication factor of these topics to 3 with kafka-reassign-partitions following an upgrade.

### Remove keystore from SMM Schema Registry client configuration if Kerberos is enabled for Schema Registry

SMM uses a Schema Registry client to fetch schemas from Schema Registry. This Schema Registry client has Kerberos authentication properties and keystore properties for mTLS. Typically, the Schema Registry server, by default, does not allow mTLS authentication. But if mTLS is enabled in the Schema Registry server, then mTLS authentication has a higher precedence than Kerberos. Therefore, the mTLS principal (from the keystore) is used for authorization with Ranger rather than the Kerberos principal. This might result in authorization failures if the mTLS principal is not added to Ranger to access the Schema Registry resources.

From now on, the Schema Registry client used by SMM does not have keystore properties for mTLS when Kerberos is enabled. As a result, even if mTLS is enabled for the Schema Registry server, the Kerberos principal is used for authentication and authorization with Ranger.

### Highly available Kafka Connect integration

SMM uses the Kafka Connect service role's REST URL to establish a connection with Connect and to serve Connect metrics. Previously, even if your Connect deployment was highly available and had multiple service roles deployed, SMM could only be configured with a single connection URL. From now on, multiple URLs can be configured. If the Connect service role that SMM is connected to fails, SMM automatically connects to a different instance that is available.

As a result of this change, the Kafka Connect Host and Kafka Connect Port properties are replaced by the Kafka Connect Rest HostPorts property. If Kafka Connect Rest HostPorts is left empty (default), SMM is automatically configured with the host, port, and protocol of the Connect service role instances belonging to the Kafka service selected with the Kafka Service SMM property.



If you previously configured Kafka Connect Host and Kafka Connect Port, the values set in the properties are automatically migrated to Kafka Connect Rest HostPorts when you upgrade.

### Partition Assignment tab [Technical Preview]

A new tab, **Assignment**, is introduced on the **Topic Details** page. This tab gives you a visual overview of the current state of the partitions and replicas of the topic. Information presented on this page includes various topic-level statistics and the replica assignment of all partitions. If rack awareness is enabled for Kafka, the replica assignment is displayed in a rack-based view. If the rack IDs follow the format of multi-level rack IDs, the rack IDs are rendered as a hierarchy.

In addition to the new tab, the **Brokers** and **Broker Details** pages now both show the rack ID of the brokers. Furthermore, a new endpoint, `/api/v1/admin/topics/{topicName}/description`, is introduced. The endpoint returns information regarding the partitions of a topic.

For more information regarding this tab as well as Kafka rack awareness, see the following resource:

- [Kafka rack awareness](#)
- [Monitoring Kafka topics](#)

### The producer of SMM Kafka interceptors can now be configured

Clients that use either of the SMM monitoring interceptors (`MonitoringConsumerInterceptor` or `MonitoringProducerInterceptor`) use a background producer to push client metrics into Kafka every 30 seconds. This background producer can from now on be configured by passing producer configurations to the clients that use the interceptor. Properties are passed to the producer with the `smm.monitoring.interceptor.producer.*` prefix.

The prefix is trimmed and the remaining part of the configuration is passed to the background producer. For example, if you want to configure the `batch.size` property for the background producer, you must set the following property:

```
smm.monitoring.interceptor.producer.batch.size
```

If you do not configure the `client.id` property of the producer, the producer uses `smm-monitoring-interceptor` as its ID instead of using an empty ID.

## What's New in Apache Hadoop YARN

Learn about the new features of YARN in Cloudera Runtime 7.1.9

### Read-only access to YARN Queue Manager UI for non-admins



**Note:** Read-only access to YARN Queue Manager UI is now a fully supported feature.

You can now allow non-admin users to access YARN Queue Manager in a read-only mode. You can either create a new user account with read-only role or use any existing user account with read-only role in Cloudera Manager to access YARN Queue Manager UI. In the read-only access mode, the user can view all the configurations but cannot make any changes to the configurations. Read-only mode will also take place during upgrades to YARN Queue Manager.

For more information, see [Provide read-only access to Queue Manager UI](#).

### Database requirement changes

After upgrading to Cloudera Data Platform (CDP) 7.1.9 or or CDP 7.1.9 cumulative hotfix (CHF) 1, you must migrate your YARN Queue Manager database to a PostgreSQL external database. This migration is needed because Queue Manager uses a different database as its backing store than what was used in previous versions of CDP 7.1.8 and below.



**Important:** If you do not want to use PostgreSQL, you may upgrade to CDP 7.1.9 CHF 2 when it becomes available as that release will no longer require you to use PostgreSQL. CDP 7.1.9 CHF 2 will allow you to continue using your embedded database.

### Fair sharing intra-queue preemption support



**Note:** Fair sharing intra-queue preemption is now a fully supported feature.

You can now enable intra-queue preemption for queues that are configured with fairness-based ordering. If the user has configured ordering-policies for queues (`yarn.scheduler.capacity.<queue-path>.ordering-policy`) to be Fair and if this new feature intra-queue preemption is enabled using YARN Queue Manager UI, then all the applications from the same user get a fair amount of resources. Thus, fair-ordering preemption now ensures that each application under a queue gets its fair-share, whether from a single user or several.

While that is for applications from a single user, for resource allocation across users, you can enforce a limit on each user's resource-usage by setting the user limits (`yarn.scheduler.capacity.<queue-path>.user-limit-factor`).

For more information about ordering policies, see [Set Ordering policies within a specific queue](#).

### Mixed resource allocation mode (Technical Preview)

You can now specify the resources in mixed types. You can specify the actual units of vcores and memory resources for each queue or specify the percentage of the total resources used by each queue or specify a weight for each queue. You can use a combination of these allocation modes. The queues under one parent can also mix their modes.

For more information, see [Mixed resource allocation mode](#).

### YARN Queue Manager Maximum Memory value update

You are now able to configure absolute capacities beyond the current cluster capacity in YARN Queue Manager. You may exceed the Maximum Memory allowed for scenarios where the nodes will be added to the cluster manually or via autoscaling.

## What's New in Apache ZooKeeper

Learn about the new features of ZooKeeper in Cloudera Runtime 7.1.9.

### Rebase ZooKeeper

CDP Private Cloud Base is updated to use Apache ZooKeeper version 3.8.1 and Apache Curator version 5.4.0 for a smoother and better functionality. Upgrade your client applications for seamless connectivity.

## Unaffected Components in this release

There are no new features for the following components in Cloudera Runtime 7.1.9.

- Apache Avro
- Apache Hadoop
- HDFS
- MapReduce
- Apache Parquet
- Apache Zeppelin

## Cloudera Runtime 7.1.9 component versions

You must be familiar with the versions of all the components in the Cloudera Runtime 7.1.9 distribution to ensure the compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

### Apache Components

The component version number has three parts, **[\*\*Apache component version number\*\*].[\*Runtime version number\*]-[\*Runtime Build number\*]**. For example, if the listed Apache HBase component version number is 2.4.6.7.1.9.0-387, 2.4.6 is the upstream Apache HBase component version, 7.1.9.0 is the Runtime version, and 387 is the Runtime build. You can also view the component version numbers in Cloudera Manager.

| Component                              | Version                                |
|----------------------------------------|----------------------------------------|
| Apache Arrow                           | 0.8.0.7.1.9.0-387                      |
| Apache Atlas                           | 2.1.0.7.1.9.0-387                      |
| Apache Calcite                         | 1.19.0.7.1.9.0-387                     |
| Apache Avro                            | 1.11.1.7.1.9.0-387                     |
| Apache Hadoop (Includes YARN and HDFS) | 3.1.1.7.1.9.0-387                      |
| Apache HBase                           | 2.4.17.7.1.9.0-387                     |
| Apache Hive                            | 3.1.3000.7.1.9.0-387                   |
| Apache Impala                          | 4.0.0.7.1.9.0-387                      |
| Apache Iceberg                         | 1.3.0.7.1.9.0-387                      |
| Apache Kafka                           | 3.4.1.7.1.9.0-387                      |
| Apache Knox                            | 1.3.0.7.1.9.0-387                      |
| Apache Kudu                            | 1.17.0.7.1.9.0-387                     |
| Apache Livy                            | 0.7.2.7.1.9.0-387                      |
| Apache MapReduce                       | 3.1.1.7.1.9.0-387                      |
| Apache Ozone                           | 1.3.0.7.1.9.0-387                      |
| Apache Oozie                           | 5.1.0.7.1.9.0-387                      |
| Apache ORC                             | 1.5.1.7.1.9.0-387                      |
| Apache Parquet                         | 1.10.99.7.1.9.0-387                    |
| Apache Phoenix                         | 5.1.1.7.1.9.0-387                      |
| Apache Ranger                          | 2.4.0.7.1.9.0-387                      |
| Apache Solr                            | 8.11.2.7.1.9.0-387                     |
| Apache Spark 2.x                       | 2.4.8.7.1.9.0-387                      |
| Apache Spark 3.x                       | 3.3.2.3.3.7190.9-1 <a href="#">CDS</a> |
| Apache Sqoop                           | 1.4.7.7.1.9.0-387                      |
| Apache Tez                             | 0.9.1.7.1.9.0-387                      |
| Apache Zeppelin                        | 0.8.2.7.1.9-387                        |
| Apache ZooKeeper                       | 3.8.1.7.1.9.0-387                      |

### Other Components

| Component      | Version             |
|----------------|---------------------|
| Cruise Control | 2.5.116.7.1.9.0-387 |

| Component                   | Version            |
|-----------------------------|--------------------|
| Data Analytics Studio       | 1.4.2.7.1.9.0-387  |
| GCS Connector               | 2.1.2.7.1.9.0-387  |
| HBase Indexer               | 1.5.0.7.5.0-387    |
| Hue                         | 4.11.0.7.1.9.0-387 |
| Search                      | 1.0.0.7.1.9.0-387  |
| Schema Registry             | 0.10.0.7.1.9.0-387 |
| Streams Messaging Manager   | 2.3.0.7.1.9.0-387  |
| Streams Replication Manager | 1.1.0.7.1.9.0-387  |

### Connectors and Encryption Components

| Component                | Version           |
|--------------------------|-------------------|
| HBase connectors         | 1.0.0.7.1.9.0-387 |
| Hive Meta Store (HMS)    | 1.0.0.7.1.9.0-387 |
| Hive on Tez              | 1.0.0.7.1.9.0-387 |
| Hive Warehouse Connector | 1.0.0.7.1.9.0-387 |
| Spark Atlas Connector    | 0.1.0.7.1.9.0-387 |
| Spark Schema Registry    | 1.1.0.7.1.9.0-387 |

## Using the Cloudera Runtime Maven repository 7.1.9

Information about using Maven to build applications with Cloudera Runtime components.

If you want to build applications or tools for use with Cloudera Runtime components and you are using Maven or Ivy for dependency management, you can pull the Cloudera Runtime artifacts from the Cloudera Maven repository. The repository is available at <https://repository.cloudera.com/artifactory/cloudera-repos/>.



**Important:** When you build an application JAR, do not include CDH JARs, because they are already provided. If you do, upgrading CDH can break your application. To avoid this situation, set the Maven dependency scope to provided. If you have already built applications which include the CDH JARs, update the dependency to set scope to provided and recompile.

The following is a sample POM (pom.xml) file:

```
<project xmlns="http://maven.apache.org/POM/4.0.0" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:schemaLocation="http://maven.apache.org/POM/4.0.0 http://maven.apache.org/maven-v4_0_0.xsd">
 <repositories>
 <repository>
 <id>cloudera</id>
 <url>https://repository.cloudera.com/artifactory/cloudera-repos/</url>
 </repository>
 </repositories>
</project>
```

### Runtime 7.1.9.0-387

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 | 3.0.0.7.1.9.0-387 |

| Project | groupId          | artifactId                   | version            |
|---------|------------------|------------------------------|--------------------|
| Atlas   | org.apache.atlas | atlas-aws-s3-bridge          | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-azure-adls-bridge      | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-classification-updater | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-client-common          | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-client-v1              | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-client-v2              | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-common                 | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-distro                 | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-docs                   | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-graphdb-api            | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-graphdb-common         | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-graphdb-janus          | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-hdfs-bridge            | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-index-repair-tool      | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-intg                   | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-janusgraph-hbase2      | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-notification           | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-plugin-classloader     | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-repository             | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-server-api             | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | atlas-testtools              | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | hbase-bridge                 | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | hbase-bridge-shim            | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | hbase-testing-util           | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | hdfs-model                   | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | hive-bridge                  | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | hive-bridge-shim             | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | impala-bridge                | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | impala-bridge-shim           | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | impala-hook-api              | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | kafka-bridge                 | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | kafka-bridge-shim            | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | navigator-to-atlas           | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | sample-app                   | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | sqoop-bridge                 | 3.0.0.7.1.9.0-387  |
| Atlas   | org.apache.atlas | sqoop-bridge-shim            | 3.0.0.7.1.9.0-387  |
| Avro    | org.apache.avro  | avro                         | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro  | avro-android                 | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro  | avro-codegen-test            | 1.11.1.7.1.9.0-387 |

| Project | groupId            | artifactId                      | version            |
|---------|--------------------|---------------------------------|--------------------|
| Avro    | org.apache.avro    | avro-compiler                   | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro    | avro-grpc                       | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro    | avro-ipc                        | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro    | avro-ipc-jetty                  | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro    | avro-ipc-netty                  | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro    | avro-mapred                     | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro    | avro-maven-plugin               | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro    | avro-perf                       | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro    | avro-protobuf                   | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro    | avro-service-archetype          | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro    | avro-test-custom-conversions    | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro    | avro-thrift                     | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro    | avro-tools                      | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro    | trevni-avro                     | 1.11.1.7.1.9.0-387 |
| Avro    | org.apache.avro    | trevni-core                     | 1.11.1.7.1.9.0-387 |
| Calcite | org.apache.calcite | calcite-babel                   | 1.19.0.7.1.9.0-387 |
| Calcite | org.apache.calcite | calcite-core                    | 1.19.0.7.1.9.0-387 |
| Calcite | org.apache.calcite | calcite-druid                   | 1.19.0.7.1.9.0-387 |
| Calcite | org.apache.calcite | calcite-linq4j                  | 1.19.0.7.1.9.0-387 |
| Calcite | org.apache.calcite | calcite-server                  | 1.19.0.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop  | hadoop-aliyun                   | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-annotations              | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-archive-logs             | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-archives                 | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-assemblies               | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-auth                     | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-aws                      | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-azure                    | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-azure-datalake           | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-build-tools              | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-client                   | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-client-api               | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-client-integration-tests | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-client-minicluster       | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-client-runtime           | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-cloud-storage            | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-common                   | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-datajoin                 | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop  | hadoop-distcp                   | 3.1.1.7.1.9.0-387  |

| Project | groupId           | artifactId                                     | version           |
|---------|-------------------|------------------------------------------------|-------------------|
| Hadoop  | org.apache.hadoop | hadoop-extras                                  | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-fs2img                                  | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-gridmix                                 | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs                                    | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-client                             | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-httpfs                             | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-native-client                      | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-nfs                                | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-rbf                                | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-kafka                                   | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-kms                                     | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-app                    | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-common                 | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-core                   | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-hs                     | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-hs-plugins             | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-jobclient              | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-nativetask             | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-shuffle                | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-uploader               | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-examples                      | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-maven-plugins                           | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-minicuster                              | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-minikdc                                 | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-nfs                                     | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-openstack                               | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-resourceestimator                       | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-rumen                                   | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-sls                                     | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-streaming                               | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-tools-dist                              | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-api                                | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-distributedshell      | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-unmanaged-am-launcher | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-client                             | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-common                             | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-registry                           | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-applicationhistoryservice   | 3.1.1.7.1.9.0-387 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-common                      | 3.1.1.7.1.9.0-387 |

| Project | groupId           | artifactId                                        | version            |
|---------|-------------------|---------------------------------------------------|--------------------|
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-nodemanager                    | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-resourcemanager                | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-router                         | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-sharedcachemanager             | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-tests                          | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timeline-pluginstorage         | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice                | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-client   | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-common   | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-server-2 | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-tests    | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-web-proxy                      | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-api                          | 3.1.1.7.1.9.0-387  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-core                         | 3.1.1.7.1.9.0-387  |
| HBase   | org.apache.hbase  | hbase-annotations                                 | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-asyncfs                                     | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-checkstyle                                  | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-client                                      | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-client-project                              | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-common                                      | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-endpoint                                    | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-examples                                    | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-external-blockcache                         | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-hadoop-compat                               | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-hadoop2-compat                              | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-hbtop                                       | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-http                                        | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-it                                          | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-logging                                     | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-mapreduce                                   | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-metrics                                     | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-metrics-api                                 | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-procedure                                   | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-protocol                                    | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-protocol-shaded                             | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-replication                                 | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-resource-bundle                             | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-rest                                        | 2.4.17.7.1.9.0-387 |
| HBase   | org.apache.hbase  | hbase-rsgroup                                     | 2.4.17.7.1.9.0-387 |



| Project | groupId          | artifactId                       | version              |
|---------|------------------|----------------------------------|----------------------|
| HBase   | org.apache.hbase | hbase-server                     | 2.4.17.7.1.9.0-387   |
| HBase   | org.apache.hbase | hbase-shaded-client              | 2.4.17.7.1.9.0-387   |
| HBase   | org.apache.hbase | hbase-shaded-client-byo-hadoop   | 2.4.17.7.1.9.0-387   |
| HBase   | org.apache.hbase | hbase-shaded-client-project      | 2.4.17.7.1.9.0-387   |
| HBase   | org.apache.hbase | hbase-shaded-mapreduce           | 2.4.17.7.1.9.0-387   |
| HBase   | org.apache.hbase | hbase-shaded-testing-util        | 2.4.17.7.1.9.0-387   |
| HBase   | org.apache.hbase | hbase-shaded-testing-util-tester | 2.4.17.7.1.9.0-387   |
| HBase   | org.apache.hbase | hbase-shell                      | 2.4.17.7.1.9.0-387   |
| HBase   | org.apache.hbase | hbase-testing-util               | 2.4.17.7.1.9.0-387   |
| HBase   | org.apache.hbase | hbase-thrift                     | 2.4.17.7.1.9.0-387   |
| HBase   | org.apache.hbase | hbase-zookeeper                  | 2.4.17.7.1.9.0-387   |
| Hive    | org.apache.hive  | hive-beeline                     | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-blobstore                   | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-classification              | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-cli                         | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-common                      | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-contrib                     | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-exec                        | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-hbase-handler               | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-hcatalog-it-unit            | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-hplsql                      | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-it-custom-serde             | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-it-minikdc                  | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-it-qfile                    | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-it-qfile-kudu               | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-it-test-serde               | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-it-unit                     | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-it-unit-hadoop2             | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-it-util                     | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-jdbc                        | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-jdbc-handler                | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-jmh                         | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-kryo-registrator            | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-kudu-handler                | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-llap-client                 | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-llap-common                 | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-llap-ext-client             | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-llap-server                 | 3.1.3000.7.1.9.0-387 |
| Hive    | org.apache.hive  | hive-llap-tez                    | 3.1.3000.7.1.9.0-387 |

| Project                  | groupId              | artifactId                               | version              |
|--------------------------|----------------------|------------------------------------------|----------------------|
| Hive                     | org.apache.hive      | hive-metastore                           | 3.1.3000.7.1.9.0-387 |
| Hive                     | org.apache.hive      | hive-pre-upgrade                         | 3.1.3000.7.1.9.0-387 |
| Hive                     | org.apache.hive      | hive-serde                               | 3.1.3000.7.1.9.0-387 |
| Hive                     | org.apache.hive      | hive-service                             | 3.1.3000.7.1.9.0-387 |
| Hive                     | org.apache.hive      | hive-service-rpc                         | 3.1.3000.7.1.9.0-387 |
| Hive                     | org.apache.hive      | hive-shims                               | 3.1.3000.7.1.9.0-387 |
| Hive                     | org.apache.hive      | hive-spark-client                        | 3.1.3000.7.1.9.0-387 |
| Hive                     | org.apache.hive      | hive-standalone-metastore                | 3.1.3000.7.1.9.0-387 |
| Hive                     | org.apache.hive      | hive-storage-api                         | 3.1.3000.7.1.9.0-387 |
| Hive                     | org.apache.hive      | hive-streaming                           | 3.1.3000.7.1.9.0-387 |
| Hive                     | org.apache.hive      | hive-testutils                           | 3.1.3000.7.1.9.0-387 |
| Hive                     | org.apache.hive      | hive-vector-code-gen                     | 3.1.3000.7.1.9.0-387 |
| Hive                     | org.apache.hive      | kafka-handler                            | 3.1.3000.7.1.9.0-387 |
| Hive Warehouse Connector | com.hortonworks.hive | hive-warehouse-connector_2.11            | 1.0.0.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | ci                                       | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | connect                                  | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | connect-api                              | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | connect-basic-auth-extension             | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | connect-cloudera-authorization-extension | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | connect-cloudera-common                  | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | connect-cloudera-secret-storage          | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | connect-cloudera-security-policies       | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | connect-file                             | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | connect-json                             | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | connect-mirror                           | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | connect-mirror-client                    | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | connect-runtime                          | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | connect-transforms                       | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | generator                                | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | kafka-clients                            | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | kafka-cloudera-metrics-reporter_2.12     | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | kafka-cloudera-metrics-reporter_2.13     | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | kafka-cloudera-plugins                   | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | kafka-examples                           | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | kafka-group-coordinator                  | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | kafka-log4j-appender                     | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | kafka-metadata                           | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | kafka-raft                               | 3.4.1.7.1.9.0-387    |
| Kafka                    | org.apache.kafka     | kafka-server-common                      | 3.4.1.7.1.9.0-387    |

| Project | groupId          | artifactId                              | version           |
|---------|------------------|-----------------------------------------|-------------------|
| Kafka   | org.apache.kafka | kafka-shell                             | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-storage                           | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-storage-api                       | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams                           | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-examples                  | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-scala_2.12                | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-scala_2.13                | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-test-utils                | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0100 | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0101 | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0102 | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0110 | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-10   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-11   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-20   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-21   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-22   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-23   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-24   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-25   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-26   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-27   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-28   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-30   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-31   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-32   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-33   | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka-tools                             | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka_2.12                              | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | kafka_2.13                              | 3.4.1.7.1.9.0-387 |
| Kafka   | org.apache.kafka | trogdor                                 | 3.4.1.7.1.9.0-387 |
| Knox    | org.apache.knox  | gateway-adapter                         | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox  | gateway-admin-ui                        | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox  | gateway-applications                    | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox  | gateway-cloud-bindings                  | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox  | gateway-demo-ldap                       | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox  | gateway-demo-ldap-launcher              | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox  | gateway-discovery-ambari                | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox  | gateway-discovery-cm                    | 1.3.0.7.1.9.0-387 |

| Project | groupId         | artifactId                                        | version           |
|---------|-----------------|---------------------------------------------------|-------------------|
| Knox    | org.apache.knox | gateway-docker                                    | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-i18n                                      | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-i18n-logging-log4j                        | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-i18n-logging-sl4j                         | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-performance-test                          | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-ha                               | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-common        | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-concat        | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-hadoop-groups | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-pseudo        | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-regex         | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-switchcase    | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-jersey                           | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-rewrite                          | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-rewrite-common                   | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-hostmap-static      | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-inbound-query-param | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-service-registry    | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-rewrite-step-encrypt-uri         | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-rewrite-step-secure-query        | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-security-authc-anon              | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-security-authz-acls              | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-security-authz-composite         | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-security-clientcert              | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-security-hadoopauth              | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-security-jwt                     | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-security-pac4j                   | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-security-preauth                 | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-security-shiro                   | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-provider-security-webappsec               | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-release                                   | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-server                                    | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-server-launcher                           | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-server-xforwarded-filter                  | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-admin                             | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-as                                | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-definitions                       | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-hashicorp-vault                   | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-hbase                             | 1.3.0.7.1.9.0-387 |

| Project | groupId         | artifactId                    | version           |
|---------|-----------------|-------------------------------|-------------------|
| Knox    | org.apache.knox | gateway-service-health        | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-hive          | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-idbroker      | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-impala        | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-jkg           | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-knoxsso       | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-knoxsout      | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-knoxtoken     | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-livy          | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-metadata      | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-nifi          | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-nifi-registry | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-remoteconfig  | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-rm            | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-session       | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-storm         | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-test          | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-tgs           | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-vault         | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-service-webhdfs       | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-shell                 | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-shell-launcher        | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-shell-release         | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-shell-samples         | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-spi                   | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-test                  | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-test-idbroker         | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-test-release-utils    | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-test-utils            | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-topology-hadoop-xml   | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-topology-simple       | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-util-common           | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-util-configinjector   | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-util-launcher         | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | gateway-util-urltemplate      | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | hadoop-examples               | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | knox-cli-launcher             | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | knox-homepage-ui              | 1.3.0.7.1.9.0-387 |
| Knox    | org.apache.knox | knox-token-generation-ui      | 1.3.0.7.1.9.0-387 |

| Project | groupId           | artifactId                  | version            |
|---------|-------------------|-----------------------------|--------------------|
| Knox    | org.apache.knox   | knox-token-management-ui    | 1.3.0.7.1.9.0-387  |
| Knox    | org.apache.knox   | webhdfs-kerb-test           | 1.3.0.7.1.9.0-387  |
| Knox    | org.apache.knox   | webhdfs-test                | 1.3.0.7.1.9.0-387  |
| Kudu    | org.apache.kudu   | kudu-backup-tools           | 1.17.0.7.1.9.0-387 |
| Kudu    | org.apache.kudu   | kudu-backup2_2.11           | 1.17.0.7.1.9.0-387 |
| Kudu    | org.apache.kudu   | kudu-backup3_2.12           | 1.17.0.7.1.9.0-387 |
| Kudu    | org.apache.kudu   | kudu-client                 | 1.17.0.7.1.9.0-387 |
| Kudu    | org.apache.kudu   | kudu-hive                   | 1.17.0.7.1.9.0-387 |
| Kudu    | org.apache.kudu   | kudu-spark2-tools_2.11      | 1.17.0.7.1.9.0-387 |
| Kudu    | org.apache.kudu   | kudu-spark2_2.11            | 1.17.0.7.1.9.0-387 |
| Kudu    | org.apache.kudu   | kudu-spark3-tools_2.12      | 1.17.0.7.1.9.0-387 |
| Kudu    | org.apache.kudu   | kudu-spark3_2.12            | 1.17.0.7.1.9.0-387 |
| Kudu    | org.apache.kudu   | kudu-test-utils             | 1.17.0.7.1.9.0-387 |
| Livy    | org.apache.livy   | livy-api                    | 0.7.2.7.1.9.0-387  |
| Livy    | org.apache.livy   | livy-client-common          | 0.7.2.7.1.9.0-387  |
| Livy    | org.apache.livy   | livy-client-http            | 0.7.2.7.1.9.0-387  |
| Livy    | org.apache.livy   | livy-core_2.11              | 0.7.2.7.1.9.0-387  |
| Livy    | org.apache.livy   | livy-examples               | 0.7.2.7.1.9.0-387  |
| Livy    | org.apache.livy   | livy-integration-test       | 0.7.2.7.1.9.0-387  |
| Livy    | org.apache.livy   | livy-repl_2.11              | 0.7.2.7.1.9.0-387  |
| Livy    | org.apache.livy   | livy-rsc                    | 0.7.2.7.1.9.0-387  |
| Livy    | org.apache.livy   | livy-scala-api_2.11         | 0.7.2.7.1.9.0-387  |
| Livy    | org.apache.livy   | livy-server                 | 0.7.2.7.1.9.0-387  |
| Livy    | org.apache.livy   | livy-test-lib               | 0.7.2.7.1.9.0-387  |
| Livy    | org.apache.livy   | livy-thriftserver           | 0.7.2.7.1.9.0-387  |
| Livy    | org.apache.livy   | livy-thriftserver-session   | 0.7.2.7.1.9.0-387  |
| Lucene  | org.apache.lucene | lucene-analyzers-common     | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-analyzers-icu        | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-analyzers-kuromoji   | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-analyzers-morfologik | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-analyzers-nori       | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-analyzers-openslp    | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-analyzers-phonetic   | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-analyzers-smartcn    | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-analyzers-stempel    | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-backward-codecs      | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-benchmark            | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-classification       | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-codecs               | 8.11.2.7.1.9.0-387 |

| Project | groupId           | artifactId                     | version            |
|---------|-------------------|--------------------------------|--------------------|
| Lucene  | org.apache.lucene | lucene-core                    | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-demo                    | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-expressions             | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-facet                   | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-grouping                | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-highlighter             | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-join                    | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-memory                  | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-misc                    | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-monitor                 | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-queries                 | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-queryparser             | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-replicator              | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-sandbox                 | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-spatial-extras          | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-spatial3d               | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-suggest                 | 8.11.2.7.1.9.0-387 |
| Lucene  | org.apache.lucene | lucene-test-framework          | 8.11.2.7.1.9.0-387 |
| Oozie   | org.apache.oozie  | oozie-client                   | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-core                     | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-distro                   | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-examples                 | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-api           | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-client        | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-server                   | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-sharelib-distcp          | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-sharelib-git             | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hcatalog        | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hive            | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hive2           | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-sharelib-oozie           | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-sharelib-spark           | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-sharelib-spark3          | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-sharelib-sqoop           | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-sharelib-streaming       | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-tools                    | 5.1.0.7.1.9.0-387  |
| Oozie   | org.apache.oozie  | oozie-zookeeper-security-tests | 5.1.0.7.1.9.0-387  |
| ORC     | org.apache.orc    | orc-core                       | 1.5.1.7.1.9.0-387  |
| ORC     | org.apache.orc    | orc-examples                   | 1.5.1.7.1.9.0-387  |

| Project | groupId          | artifactId                    | version           |
|---------|------------------|-------------------------------|-------------------|
| ORC     | org.apache.orc   | orc-mapreduce                 | 1.5.1.7.1.9.0-387 |
| ORC     | org.apache.orc   | orc-shims                     | 1.5.1.7.1.9.0-387 |
| ORC     | org.apache.orc   | orc-tools                     | 1.5.1.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-annotation-processing    | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-client                   | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-common                   | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-config                   | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-container-service        | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-docs                     | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-erasurecode              | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-client | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-interface-admin          | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-interface-client         | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-interface-server         | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-managed-rocksdb          | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-rocks-native             | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-server-framework         | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-server-scm               | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-test-utils               | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | hdds-tools                    | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | mini-chaos-tests              | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-client                  | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-common                  | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-csi                     | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-datanode                | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-dist                    | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-filesystem              | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-filesystem-common       | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-filesystem-hadoop2      | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-filesystem-hadoop3      | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-filesystem-shaded       | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-httpfgateway            | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-insight                 | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-integration-test        | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-interface-client        | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-interface-storage       | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-manager                 | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-network-tests           | 1.2.0.7.1.9.0-387 |
| Ozone   | org.apache.ozone | ozone-recon                   | 1.2.0.7.1.9.0-387 |



| Project | groupId            | artifactId                             | version             |
|---------|--------------------|----------------------------------------|---------------------|
| Ozone   | org.apache.ozone   | ozone-reconcodegen                     | 1.2.0.7.1.9.0-387   |
| Ozone   | org.apache.ozone   | ozone-s3-secret-store                  | 1.2.0.7.1.9.0-387   |
| Ozone   | org.apache.ozone   | ozone-s3gateway                        | 1.2.0.7.1.9.0-387   |
| Ozone   | org.apache.ozone   | ozone-tools                            | 1.2.0.7.1.9.0-387   |
| Ozone   | org.apache.ozone   | rocksdb-checkpoint-differ              | 1.2.0.7.1.9.0-387   |
| Parquet | org.apache.parquet | parquet-avro                           | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-cascading                      | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-cascading3                     | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-column                         | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-common                         | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-encoding                       | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-format-structures              | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-generator                      | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-hadoop                         | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-hadoop-bundle                  | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-jackson                        | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-pig                            | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-pig-bundle                     | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-protobuf                       | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-scala_2.10                     | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-thrift                         | 1.10.99.7.1.9.0-387 |
| Parquet | org.apache.parquet | parquet-tools                          | 1.10.99.7.1.9.0-387 |
| Phoenix | org.apache.phoenix | phoenix-client-embedded-hbase-2.4      | 5.1.1.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-client-hbase-2.4               | 5.1.1.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-connectors-phoenix5-compatible | 6.0.0.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-core                           | 5.1.1.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.1.6         | 5.1.1.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.2.5         | 5.1.1.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.3.0         | 5.1.1.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.4.0         | 5.1.1.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.4.1         | 5.1.1.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-pherf                          | 5.1.1.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-queryserver                    | 6.0.0.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-client             | 6.0.0.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-it                 | 6.0.0.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-load-balancer      | 6.0.0.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-orchestrator       | 6.0.0.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-server-hbase-2.4               | 5.1.1.7.1.9.0-387   |
| Phoenix | org.apache.phoenix | phoenix-tracing-webapp                 | 5.1.1.7.1.9.0-387   |

| Project | groupId            | artifactId                  | version           |
|---------|--------------------|-----------------------------|-------------------|
| Phoenix | org.apache.phoenix | phoenix5-hive               | 6.0.0.7.1.9.0-387 |
| Phoenix | org.apache.phoenix | phoenix5-hive-shaded        | 6.0.0.7.1.9.0-387 |
| Phoenix | org.apache.phoenix | phoenix5-spark              | 6.0.0.7.1.9.0-387 |
| Phoenix | org.apache.phoenix | phoenix5-spark-shaded       | 6.0.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | conditions-enrichers        | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | credentialbuilder           | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | embeddedwebserver           | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | jisql                       | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ldapconfigcheck             | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-adls-plugin          | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-atlas-plugin         | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-atlas-plugin-shim    | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-authn                | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-common-ha            | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-distro               | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-examples-distro      | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-hbase-plugin         | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-hbase-plugin-shim    | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-hdfs-plugin          | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-hdfs-plugin-shim     | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-hive-plugin          | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-hive-plugin-shim     | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-intg                 | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-kafka-connect-plugin | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-kafka-plugin         | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-kafka-plugin-shim    | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-kms                  | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-kms-plugin           | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-kms-plugin-shim      | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-knox-plugin          | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-knox-plugin-shim     | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-kudu-plugin          | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-kylin-plugin         | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-kylin-plugin-shim    | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-metrics              | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-nifi-plugin          | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-nifi-registry-plugin | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-ozone-plugin         | 2.4.0.7.1.9.0-387 |
| Ranger  | org.apache.ranger  | ranger-ozone-plugin-shim    | 2.4.0.7.1.9.0-387 |

| Project | groupId           | artifactId                    | version            |
|---------|-------------------|-------------------------------|--------------------|
| Ranger  | org.apache.ranger | ranger-plugin-classloader     | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-plugins-audit          | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-plugins-common         | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-plugins-cred           | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-plugins-installer      | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-policymigration        | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-raz-adls               | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-raz-chained-plugins    | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-raz-hook-abfs          | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-raz-intg               | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-raz-processor          | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-rms-common             | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-rms-hive               | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-rms-plugins-common     | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-rms-webapp             | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-sampleapp-plugin       | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-schema-registry-plugin | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-solr-plugin            | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-solr-plugin-shim       | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin           | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin-shim      | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-storm-plugin           | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-storm-plugin-shim      | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-tagsync                | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-tools                  | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-util                   | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-yarn-plugin            | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ranger-yarn-plugin-shim       | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | sample-client                 | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | sampleapp                     | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | shaded-raz-hook-abfs          | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | ugsync-util                   | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | unixauthclient                | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | unixauthservice               | 2.4.0.7.1.9.0-387  |
| Ranger  | org.apache.ranger | unixusersync                  | 2.4.0.7.1.9.0-387  |
| Solr    | org.apache.solr   | solr-analysis-extras          | 8.11.2.7.1.9.0-387 |
| Solr    | org.apache.solr   | solr-analytics                | 8.11.2.7.1.9.0-387 |
| Solr    | org.apache.solr   | solr-cell                     | 8.11.2.7.1.9.0-387 |
| Solr    | org.apache.solr   | solr-core                     | 8.11.2.7.1.9.0-387 |

| Project | groupId          | artifactId                               | version            |
|---------|------------------|------------------------------------------|--------------------|
| Solr    | org.apache.solr  | solr-dataimporthandler                   | 8.11.2.7.1.9.0-387 |
| Solr    | org.apache.solr  | solr-dataimporthandler-extras            | 8.11.2.7.1.9.0-387 |
| Solr    | org.apache.solr  | solr-gcs-repository                      | 8.11.2.7.1.9.0-387 |
| Solr    | org.apache.solr  | solr-jaegertracer-configurator           | 8.11.2.7.1.9.0-387 |
| Solr    | org.apache.solr  | solr-langid                              | 8.11.2.7.1.9.0-387 |
| Solr    | org.apache.solr  | solr-ltr                                 | 8.11.2.7.1.9.0-387 |
| Solr    | org.apache.solr  | solr-prometheus-exporter                 | 8.11.2.7.1.9.0-387 |
| Solr    | org.apache.solr  | solr-s3-repository                       | 8.11.2.7.1.9.0-387 |
| Solr    | org.apache.solr  | solr-security-util                       | 8.11.2.7.1.9.0-387 |
| Solr    | org.apache.solr  | solr-solrj                               | 8.11.2.7.1.9.0-387 |
| Solr    | org.apache.solr  | solr-test-framework                      | 8.11.2.7.1.9.0-387 |
| Solr    | org.apache.solr  | solr-velocity                            | 8.11.2.7.1.9.0-387 |
| Spark   | org.apache.spark | spark-avro_2.11                          | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-catalyst_2.11                      | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-core_2.11                          | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-graphx_2.11                        | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-hadoop-cloud_2.11                  | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-hive_2.11                          | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-kubernetes_2.11                    | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-kvstore_2.11                       | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-launcher_2.11                      | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-mllib-local_2.11                   | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-mllib_2.11                         | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-network-common_2.11                | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-network-shuffle_2.11               | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-network-yarn_2.11                  | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-repl_2.11                          | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-sketch_2.11                        | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-sql-kafka-0-10_2.11                | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-sql_2.11                           | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-streaming-kafka-0-10-assembly_2.11 | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-streaming-kafka-0-10_2.11          | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-streaming_2.11                     | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-tags_2.11                          | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-token-provider-kafka-0-10_2.11     | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-unsafe_2.11                        | 2.4.8.7.1.9.0-387  |
| Spark   | org.apache.spark | spark-yarn_2.11                          | 2.4.8.7.1.9.0-387  |
| Sqoop   | org.apache.sqoop | sqoop                                    | 1.4.7.7.1.9.0-387  |
| Sqoop   | org.apache.sqoop | sqoop-test                               | 1.4.7.7.1.9.0-387  |

| Project   | groupId              | artifactId                          | version           |
|-----------|----------------------|-------------------------------------|-------------------|
| Tez       | org.apache.tez       | hadoop-shim                         | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | hadoop-shim-2.8                     | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-api                             | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-aux-services                    | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-common                          | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-dag                             | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-examples                        | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-ext-service-tests               | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-history-parser                  | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-javadoc-tools                   | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-job-analyzer                    | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-mapreduce                       | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-protobuf-history-plugin         | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-runtime-internals               | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-runtime-library                 | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-tests                           | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-yarn-timeline-cache-plugin      | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history           | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history-with-acls | 0.9.1.7.1.9.0-387 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history-with-fs   | 0.9.1.7.1.9.0-387 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-angular                    | 0.8.2.7.1.9.0-387 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-display                    | 0.8.2.7.1.9.0-387 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-interpreter                | 0.8.2.7.1.9.0-387 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-jdbc                       | 0.8.2.7.1.9.0-387 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-jupyter                    | 0.8.2.7.1.9.0-387 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-livy                       | 0.8.2.7.1.9.0-387 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-markdown                   | 0.8.2.7.1.9.0-387 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-server                     | 0.8.2.7.1.9.0-387 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-shell                      | 0.8.2.7.1.9.0-387 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-zengine                    | 0.8.2.7.1.9.0-387 |
| ZooKeeper | org.apache.zookeeper | zookeeper                           | 3.8.1.7.1.9.0-387 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-fatjar            | 3.8.1.7.1.9.0-387 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-loggraph          | 3.8.1.7.1.9.0-387 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-rest              | 3.8.1.7.1.9.0-387 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-zooinspector      | 3.8.1.7.1.9.0-387 |
| ZooKeeper | org.apache.zookeeper | zookeeper-it                        | 3.8.1.7.1.9.0-387 |
| ZooKeeper | org.apache.zookeeper | zookeeper-jute                      | 3.8.1.7.1.9.0-387 |
| ZooKeeper | org.apache.zookeeper | zookeeper-prometheus-metrics        | 3.8.1.7.1.9.0-387 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-election          | 3.8.1.7.1.9.0-387 |

| Project   | groupId              | artifactId              | version           |
|-----------|----------------------|-------------------------|-------------------|
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-lock  | 3.8.1.7.1.9.0-387 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-queue | 3.8.1.7.1.9.0-387 |

## Runtime 7.1.9.2-10

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          | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-aws-s3-bridge          | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-azure-adls-bridge      | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-classification-updater | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-client-common          | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-client-v1              | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-client-v2              | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-common                 | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-distro                 | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-docs                   | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-graphdb-api            | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-graphdb-common         | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-graphdb-janus          | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-hdfs-bridge            | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-index-repair-tool      | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-intg                   | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-janusgraph-hbase2      | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-notification           | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-plugin-classloader     | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-repository             | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-server-api             | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | atlas-testtools              | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | hbase-bridge                 | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | hbase-bridge-shim            | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | hbase-testing-util           | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | hdfs-model                   | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | hive-bridge                  | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | hive-bridge-shim             | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | impala-bridge                | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | impala-bridge-shim           | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | impala-hook-api              | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | kafka-bridge                 | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | kafka-bridge-shim            | 3.0.0.7.1.9.2-10 |
| Atlas   | org.apache.atlas | navigator-to-atlas           | 3.0.0.7.1.9.2-10 |

| Project | groupId            | artifactId                      | version           |
|---------|--------------------|---------------------------------|-------------------|
| Atlas   | org.apache.atlas   | sample-app                      | 3.0.0.7.1.9.2-10  |
| Atlas   | org.apache.atlas   | sqoop-bridge                    | 3.0.0.7.1.9.2-10  |
| Atlas   | org.apache.atlas   | sqoop-bridge-shim               | 3.0.0.7.1.9.2-10  |
| Avro    | org.apache.avro    | avro                            | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-android                    | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-codegen-test               | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-compiler                   | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-grpc                       | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-ipc                        | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-ipc-jetty                  | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-ipc-netty                  | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-mapred                     | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-maven-plugin               | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-perf                       | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-protobuf                   | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-service-archetype          | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-test-custom-conversions    | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-thrift                     | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | avro-tools                      | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | trevni-avro                     | 1.11.1.7.1.9.2-10 |
| Avro    | org.apache.avro    | trevni-core                     | 1.11.1.7.1.9.2-10 |
| Calcite | org.apache.calcite | calcite-babel                   | 1.19.0.7.1.9.2-10 |
| Calcite | org.apache.calcite | calcite-core                    | 1.19.0.7.1.9.2-10 |
| Calcite | org.apache.calcite | calcite-druid                   | 1.19.0.7.1.9.2-10 |
| Calcite | org.apache.calcite | calcite-linq4j                  | 1.19.0.7.1.9.2-10 |
| Calcite | org.apache.calcite | calcite-server                  | 1.19.0.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop  | hadoop-aliyun                   | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop  | hadoop-annotations              | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop  | hadoop-archive-logs             | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop  | hadoop-archives                 | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop  | hadoop-assemblies               | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop  | hadoop-auth                     | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop  | hadoop-aws                      | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop  | hadoop-azure                    | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop  | hadoop-azure-datalake           | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop  | hadoop-build-tools              | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop  | hadoop-client                   | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop  | hadoop-client-api               | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop  | hadoop-client-integration-tests | 3.1.1.7.1.9.2-10  |

| Project | groupId           | artifactId                                | version          |
|---------|-------------------|-------------------------------------------|------------------|
| Hadoop  | org.apache.hadoop | hadoop-client-minicluster                 | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-client-runtime                     | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-cloud-storage                      | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-common                             | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-datajoin                           | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-distcp                             | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-extras                             | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-fs2img                             | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-gridmix                            | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs                               | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-client                        | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-httpfs                        | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-native-client                 | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-nfs                           | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-rbf                           | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-kafka                              | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-kms                                | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-app               | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-common            | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-core              | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-hs                | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-hs-plugins        | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-jobclient         | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-nativetask        | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-shuffle           | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-uploader          | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-examples                 | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-maven-plugins                      | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-minicluster                        | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-minikdc                            | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-nfs                                | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-openstack                          | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-resourceestimator                  | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-rumen                              | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-sls                                | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-streaming                          | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-tools-dist                         | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-api                           | 3.1.1.7.1.9.2-10 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-distributedshell | 3.1.1.7.1.9.2-10 |



| Project | groupId           | artifactId                                        | version           |
|---------|-------------------|---------------------------------------------------|-------------------|
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-unmanaged-am-launcher    | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-client                                | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-common                                | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-registry                              | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-applicationhistoryservice      | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-common                         | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-nodemanager                    | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-resourcemanager                | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-router                         | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-sharedcachemanager             | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-tests                          | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timeline-pluginstorage         | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice                | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-client   | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-common   | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-server-2 | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-tests    | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-web-proxy                      | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-api                          | 3.1.1.7.1.9.2-10  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-core                         | 3.1.1.7.1.9.2-10  |
| HBase   | org.apache.hbase  | hbase-annotations                                 | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-asyncfs                                     | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-checkstyle                                  | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-client                                      | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-client-project                              | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-common                                      | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-endpoint                                    | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-examples                                    | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-external-blockcache                         | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-hadoop-compat                               | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-hadoop2-compat                              | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-hbtop                                       | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-http                                        | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-it                                          | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-logging                                     | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-mapreduce                                   | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-metrics                                     | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-metrics-api                                 | 2.4.17.7.1.9.2-10 |
| HBase   | org.apache.hbase  | hbase-procedure                                   | 2.4.17.7.1.9.2-10 |

| Project | groupId          | artifactId                       | version             |
|---------|------------------|----------------------------------|---------------------|
| HBase   | org.apache.hbase | hbase-protocol                   | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-protocol-shaded            | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-replication                | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-resource-bundle            | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-rest                       | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-rsgroup                    | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-server                     | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-shaded-client              | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-shaded-client-byo-hadoop   | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-shaded-client-project      | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-shaded-mapreduce           | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-shaded-testing-util        | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-shaded-testing-util-tester | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-shell                      | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-testing-util               | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-thrift                     | 2.4.17.7.1.9.2-10   |
| HBase   | org.apache.hbase | hbase-zookeeper                  | 2.4.17.7.1.9.2-10   |
| Hive    | org.apache.hive  | hive-beeline                     | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-blobstore                   | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-classification              | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-cli                         | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-common                      | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-contrib                     | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-exec                        | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-hbase-handler               | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-hcatalog-it-unit            | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-hplsql                      | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-it-custom-serde             | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-it-minikdc                  | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-it-qfile                    | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-it-qfile-kudu               | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-it-test-serde               | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-it-unit                     | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-it-unit-hadoop2             | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-it-util                     | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-jdbc                        | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-jdbc-handler                | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-jmh                         | 3.1.3000.7.1.9.2-10 |
| Hive    | org.apache.hive  | hive-kryo-registrator            | 3.1.3000.7.1.9.2-10 |

| Project                  | groupId              | artifactId                               | version             |
|--------------------------|----------------------|------------------------------------------|---------------------|
| Hive                     | org.apache.hive      | hive-kudu-handler                        | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-llap-client                         | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-llap-common                         | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-llap-ext-client                     | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-llap-server                         | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-llap-tez                            | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-metastore                           | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-pre-upgrade                         | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-serde                               | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-service                             | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-service-rpc                         | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-shims                               | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-spark-client                        | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-standalone-metastore                | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-storage-api                         | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-streaming                           | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-testutils                           | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | hive-vector-code-gen                     | 3.1.3000.7.1.9.2-10 |
| Hive                     | org.apache.hive      | kafka-handler                            | 3.1.3000.7.1.9.2-10 |
| Hive Warehouse Connector | com.hortonworks.hive | hive-warehouse-connector_2.11            | 1.0.0.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | ci                                       | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | connect                                  | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | connect-api                              | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | connect-basic-auth-extension             | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | connect-cloudera-authorization-extension | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | connect-cloudera-common                  | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | connect-cloudera-secret-storage          | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | connect-cloudera-security-policies       | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | connect-file                             | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | connect-json                             | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | connect-mirror                           | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | connect-mirror-client                    | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | connect-runtime                          | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | connect-transforms                       | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | generator                                | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | kafka-clients                            | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | kafka-cloudera-metrics-reporter_2.12     | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | kafka-cloudera-metrics-reporter_2.13     | 3.4.1.7.1.9.2-10    |
| Kafka                    | org.apache.kafka     | kafka-cloudera-plugins                   | 3.4.1.7.1.9.2-10    |

| Project | groupId          | artifactId                              | version          |
|---------|------------------|-----------------------------------------|------------------|
| Kafka   | org.apache.kafka | kafka-examples                          | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-group-coordinator                 | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-log4j-appender                    | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-metadata                          | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-raft                              | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-server-common                     | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-shell                             | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-storage                           | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-storage-api                       | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams                           | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-examples                  | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-scala_2.12                | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-scala_2.13                | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-test-utils                | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0100 | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0101 | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0102 | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0110 | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-10   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-11   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-20   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-21   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-22   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-23   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-24   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-25   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-26   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-27   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-28   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-30   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-31   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-32   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-33   | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka-tools                             | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka_2.12                              | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | kafka_2.13                              | 3.4.1.7.1.9.2-10 |
| Kafka   | org.apache.kafka | trogdor                                 | 3.4.1.7.1.9.2-10 |
| Knox    | org.apache.knox  | gateway-adapter                         | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox  | gateway-admin-ui                        | 1.3.0.7.1.9.2-10 |

| Project | groupId         | artifactId                                        | version          |
|---------|-----------------|---------------------------------------------------|------------------|
| Knox    | org.apache.knox | gateway-applications                              | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-cloud-bindings                            | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-demo-ldap                                 | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-demo-ldap-launcher                        | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-discovery-ambari                          | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-discovery-cm                              | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-docker                                    | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-i18n                                      | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-i18n-logging-log4j                        | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-i18n-logging-sl4j                         | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-performance-test                          | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-ha                               | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-common        | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-concat        | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-hadoop-groups | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-pseudo        | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-regex         | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-switchcase    | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-jersey                           | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-rewrite                          | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-rewrite-common                   | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-hostmap-static      | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-inbound-query-param | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-service-registry    | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-rewrite-step-encrypt-uri         | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-rewrite-step-secure-query        | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-security-authc-anon              | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-security-authz-acls              | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-security-authz-composite         | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-security-clientcert              | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-security-hadoopauth              | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-security-jwt                     | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-security-pac4j                   | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-security-preauth                 | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-security-shiro                   | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-provider-security-webappsec               | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-release                                   | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-server                                    | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-server-launcher                           | 1.3.0.7.1.9.2-10 |

| Project | groupId         | artifactId                       | version          |
|---------|-----------------|----------------------------------|------------------|
| Knox    | org.apache.knox | gateway-server-xforwarded-filter | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-admin            | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-as               | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-definitions      | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-hashicorp-vault  | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-hbase            | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-health           | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-hive             | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-idbroker         | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-impala           | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-jkg              | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-knoxsso          | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-knoxssout        | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-knoxtoken        | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-livy             | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-metadata         | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-nifi             | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-nifi-registry    | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-remoteconfig     | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-rm               | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-session          | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-storm            | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-test             | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-tgs              | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-vault            | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-service-webhdfs          | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-shell                    | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-shell-launcher           | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-shell-release            | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-shell-samples            | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-spi                      | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-test                     | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-test-idbroker            | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-test-release-utils       | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-test-utils               | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-topology-hadoop-xml      | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-topology-simple          | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-util-common              | 1.3.0.7.1.9.2-10 |
| Knox    | org.apache.knox | gateway-util-configinjector      | 1.3.0.7.1.9.2-10 |

| Project | groupId           | artifactId                  | version           |
|---------|-------------------|-----------------------------|-------------------|
| Knox    | org.apache.knox   | gateway-util-launcher       | 1.3.0.7.1.9.2-10  |
| Knox    | org.apache.knox   | gateway-util-urltemplate    | 1.3.0.7.1.9.2-10  |
| Knox    | org.apache.knox   | hadoop-examples             | 1.3.0.7.1.9.2-10  |
| Knox    | org.apache.knox   | knox-cli-launcher           | 1.3.0.7.1.9.2-10  |
| Knox    | org.apache.knox   | knox-homepage-ui            | 1.3.0.7.1.9.2-10  |
| Knox    | org.apache.knox   | knox-token-generation-ui    | 1.3.0.7.1.9.2-10  |
| Knox    | org.apache.knox   | knox-token-management-ui    | 1.3.0.7.1.9.2-10  |
| Knox    | org.apache.knox   | webhdfs-kerb-test           | 1.3.0.7.1.9.2-10  |
| Knox    | org.apache.knox   | webhdfs-test                | 1.3.0.7.1.9.2-10  |
| Kudu    | org.apache.kudu   | kudu-backup-tools           | 1.17.0.7.1.9.2-10 |
| Kudu    | org.apache.kudu   | kudu-backup2_2.11           | 1.17.0.7.1.9.2-10 |
| Kudu    | org.apache.kudu   | kudu-backup3_2.12           | 1.17.0.7.1.9.2-10 |
| Kudu    | org.apache.kudu   | kudu-client                 | 1.17.0.7.1.9.2-10 |
| Kudu    | org.apache.kudu   | kudu-hive                   | 1.17.0.7.1.9.2-10 |
| Kudu    | org.apache.kudu   | kudu-spark2-tools_2.11      | 1.17.0.7.1.9.2-10 |
| Kudu    | org.apache.kudu   | kudu-spark2_2.11            | 1.17.0.7.1.9.2-10 |
| Kudu    | org.apache.kudu   | kudu-spark3-tools_2.12      | 1.17.0.7.1.9.2-10 |
| Kudu    | org.apache.kudu   | kudu-spark3_2.12            | 1.17.0.7.1.9.2-10 |
| Kudu    | org.apache.kudu   | kudu-test-utils             | 1.17.0.7.1.9.2-10 |
| Livy    | org.apache.livy   | livy-api                    | 0.7.2.7.1.9.2-10  |
| Livy    | org.apache.livy   | livy-client-common          | 0.7.2.7.1.9.2-10  |
| Livy    | org.apache.livy   | livy-client-http            | 0.7.2.7.1.9.2-10  |
| Livy    | org.apache.livy   | livy-core_2.11              | 0.7.2.7.1.9.2-10  |
| Livy    | org.apache.livy   | livy-examples               | 0.7.2.7.1.9.2-10  |
| Livy    | org.apache.livy   | livy-integration-test       | 0.7.2.7.1.9.2-10  |
| Livy    | org.apache.livy   | livy-repl_2.11              | 0.7.2.7.1.9.2-10  |
| Livy    | org.apache.livy   | livy-rsc                    | 0.7.2.7.1.9.2-10  |
| Livy    | org.apache.livy   | livy-scala-api_2.11         | 0.7.2.7.1.9.2-10  |
| Livy    | org.apache.livy   | livy-server                 | 0.7.2.7.1.9.2-10  |
| Livy    | org.apache.livy   | livy-test-lib               | 0.7.2.7.1.9.2-10  |
| Livy    | org.apache.livy   | livy-thriftserver           | 0.7.2.7.1.9.2-10  |
| Livy    | org.apache.livy   | livy-thriftserver-session   | 0.7.2.7.1.9.2-10  |
| Lucene  | org.apache.lucene | lucene-analyzers-common     | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-analyzers-icu        | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-analyzers-kuromoji   | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-analyzers-morfologik | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-analyzers-nori       | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-analyzers-openssl    | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-analyzers-phonetic   | 8.11.2.7.1.9.2-10 |

| Project | groupId           | artifactId               | version           |
|---------|-------------------|--------------------------|-------------------|
| Lucene  | org.apache.lucene | lucene-analyzers-smartcn | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-analyzers-stempel | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-backward-codecs   | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-benchmark         | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-classification    | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-codecs            | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-core              | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-demo              | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-expressions       | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-facet             | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-grouping          | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-highlighter       | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-join              | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-memory            | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-misc              | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-monitor           | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-queries           | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-queryparser       | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-replicator        | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-sandbox           | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-spatial-extras    | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-spatial3d         | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-suggest           | 8.11.2.7.1.9.2-10 |
| Lucene  | org.apache.lucene | lucene-test-framework    | 8.11.2.7.1.9.2-10 |
| Oozie   | org.apache.oozie  | oozie-client             | 5.1.0.7.1.9.2-10  |
| Oozie   | org.apache.oozie  | oozie-core               | 5.1.0.7.1.9.2-10  |
| Oozie   | org.apache.oozie  | oozie-distro             | 5.1.0.7.1.9.2-10  |
| Oozie   | org.apache.oozie  | oozie-examples           | 5.1.0.7.1.9.2-10  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-api     | 5.1.0.7.1.9.2-10  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-client  | 5.1.0.7.1.9.2-10  |
| Oozie   | org.apache.oozie  | oozie-server             | 5.1.0.7.1.9.2-10  |
| Oozie   | org.apache.oozie  | oozie-sharelib-distcp    | 5.1.0.7.1.9.2-10  |
| Oozie   | org.apache.oozie  | oozie-sharelib-git       | 5.1.0.7.1.9.2-10  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hcatalog  | 5.1.0.7.1.9.2-10  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hive      | 5.1.0.7.1.9.2-10  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hive2     | 5.1.0.7.1.9.2-10  |
| Oozie   | org.apache.oozie  | oozie-sharelib-oozie     | 5.1.0.7.1.9.2-10  |
| Oozie   | org.apache.oozie  | oozie-sharelib-spark     | 5.1.0.7.1.9.2-10  |
| Oozie   | org.apache.oozie  | oozie-sharelib-spark3    | 5.1.0.7.1.9.2-10  |



| Project | groupId          | artifactId                     | version          |
|---------|------------------|--------------------------------|------------------|
| Oozie   | org.apache.oozie | oozie-sharelib-sqoop           | 5.1.0.7.1.9.2-10 |
| Oozie   | org.apache.oozie | oozie-sharelib-streaming       | 5.1.0.7.1.9.2-10 |
| Oozie   | org.apache.oozie | oozie-tools                    | 5.1.0.7.1.9.2-10 |
| Oozie   | org.apache.oozie | oozie-zookeeper-security-tests | 5.1.0.7.1.9.2-10 |
| ORC     | org.apache.orc   | orc-core                       | 1.5.1.7.1.9.2-10 |
| ORC     | org.apache.orc   | orc-examples                   | 1.5.1.7.1.9.2-10 |
| ORC     | org.apache.orc   | orc-mapreduce                  | 1.5.1.7.1.9.2-10 |
| ORC     | org.apache.orc   | orc-shims                      | 1.5.1.7.1.9.2-10 |
| ORC     | org.apache.orc   | orc-tools                      | 1.5.1.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-annotation-processing     | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-client                    | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-common                    | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-config                    | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-container-service         | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-docs                      | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-erasurecode               | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-client  | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-server  | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-test    | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-interface-admin           | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-interface-client          | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-interface-server          | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-managed-rocksdbs          | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-rocks-native              | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-server-framework          | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-server-scm                | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-test-utils                | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | hdds-tools                     | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | mini-chaos-tests               | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | ozone-client                   | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | ozone-common                   | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | ozone-csi                      | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | ozone-datanode                 | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | ozone-dist                     | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | ozone-filesystem               | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | ozone-filesystem-common        | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | ozone-filesystem-hadoop2       | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | ozone-filesystem-hadoop3       | 1.2.0.7.1.9.2-10 |
| Ozone   | org.apache.ozone | ozone-filesystem-shaded        | 1.2.0.7.1.9.2-10 |

| Project | groupId            | artifactId                             | version            |
|---------|--------------------|----------------------------------------|--------------------|
| Ozone   | org.apache.ozone   | ozone-httpfsgateway                    | 1.2.0.7.1.9.2-10   |
| Ozone   | org.apache.ozone   | ozone-insight                          | 1.2.0.7.1.9.2-10   |
| Ozone   | org.apache.ozone   | ozone-integration-test                 | 1.2.0.7.1.9.2-10   |
| Ozone   | org.apache.ozone   | ozone-interface-client                 | 1.2.0.7.1.9.2-10   |
| Ozone   | org.apache.ozone   | ozone-interface-storage                | 1.2.0.7.1.9.2-10   |
| Ozone   | org.apache.ozone   | ozone-manager                          | 1.2.0.7.1.9.2-10   |
| Ozone   | org.apache.ozone   | ozone-network-tests                    | 1.2.0.7.1.9.2-10   |
| Ozone   | org.apache.ozone   | ozone-recon                            | 1.2.0.7.1.9.2-10   |
| Ozone   | org.apache.ozone   | ozone-reconcodegen                     | 1.2.0.7.1.9.2-10   |
| Ozone   | org.apache.ozone   | ozone-s3-secret-store                  | 1.2.0.7.1.9.2-10   |
| Ozone   | org.apache.ozone   | ozone-s3gateway                        | 1.2.0.7.1.9.2-10   |
| Ozone   | org.apache.ozone   | ozone-tools                            | 1.2.0.7.1.9.2-10   |
| Ozone   | org.apache.ozone   | rocksdb-checkpoint-differ              | 1.2.0.7.1.9.2-10   |
| Parquet | org.apache.parquet | parquet-avro                           | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-cascading                      | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-cascading3                     | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-column                         | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-common                         | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-encoding                       | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-format-structures              | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-generator                      | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-hadoop                         | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-hadoop-bundle                  | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-jackson                        | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-pig                            | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-pig-bundle                     | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-protobuf                       | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-scala_2.10                     | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-thrift                         | 1.10.99.7.1.9.2-10 |
| Parquet | org.apache.parquet | parquet-tools                          | 1.10.99.7.1.9.2-10 |
| Phoenix | org.apache.phoenix | phoenix-client-embedded-hbase-2.4      | 5.1.1.7.1.9.2-10   |
| Phoenix | org.apache.phoenix | phoenix-client-hbase-2.4               | 5.1.1.7.1.9.2-10   |
| Phoenix | org.apache.phoenix | phoenix-connectors-phoenix5-compatible | 6.0.0.7.1.9.2-10   |
| Phoenix | org.apache.phoenix | phoenix-core                           | 5.1.1.7.1.9.2-10   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.1.6         | 5.1.1.7.1.9.2-10   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.2.5         | 5.1.1.7.1.9.2-10   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.3.0         | 5.1.1.7.1.9.2-10   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.4.0         | 5.1.1.7.1.9.2-10   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.4.1         | 5.1.1.7.1.9.2-10   |

| Project | groupId            | artifactId                        | version          |
|---------|--------------------|-----------------------------------|------------------|
| Phoenix | org.apache.phoenix | phoenix-pherf                     | 5.1.1.7.1.9.2-10 |
| Phoenix | org.apache.phoenix | phoenix-queryserver               | 6.0.0.7.1.9.2-10 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-client        | 6.0.0.7.1.9.2-10 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-it            | 6.0.0.7.1.9.2-10 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-load-balancer | 6.0.0.7.1.9.2-10 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-orchestrator  | 6.0.0.7.1.9.2-10 |
| Phoenix | org.apache.phoenix | phoenix-server-hbase-2.4          | 5.1.1.7.1.9.2-10 |
| Phoenix | org.apache.phoenix | phoenix-tracing-webapp            | 5.1.1.7.1.9.2-10 |
| Phoenix | org.apache.phoenix | phoenix5-hive                     | 6.0.0.7.1.9.2-10 |
| Phoenix | org.apache.phoenix | phoenix5-hive-shaded              | 6.0.0.7.1.9.2-10 |
| Phoenix | org.apache.phoenix | phoenix5-spark                    | 6.0.0.7.1.9.2-10 |
| Phoenix | org.apache.phoenix | phoenix5-spark-shaded             | 6.0.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | conditions-enrichers              | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | credentialbuilder                 | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | embeddedwebserver                 | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | jisql                             | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ldapconfigcheck                   | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-adls-plugin                | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-atlas-plugin               | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-atlas-plugin-shim          | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-authn                      | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-common-ha                  | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-distro                     | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-examples-distro            | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-hbase-plugin               | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-hbase-plugin-shim          | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-hdfs-plugin                | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-hdfs-plugin-shim           | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-hive-plugin                | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-hive-plugin-shim           | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-intg                       | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-kafka-connect-plugin       | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-kafka-plugin               | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-kafka-plugin-shim          | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-kms                        | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-kms-plugin                 | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-kms-plugin-shim            | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-knox-plugin                | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger  | ranger-knox-plugin-shim           | 2.4.0.7.1.9.2-10 |

| Project | groupId           | artifactId                    | version          |
|---------|-------------------|-------------------------------|------------------|
| Ranger  | org.apache.ranger | ranger-kudu-plugin            | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-kylin-plugin           | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-kylin-plugin-shim      | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-metrics                | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-nifi-plugin            | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-nifi-registry-plugin   | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-ozone-plugin           | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-ozone-plugin-shim      | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-plugin-classloader     | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-plugins-audit          | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-plugins-common         | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-plugins-cred           | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-plugins-installer      | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-policymigration        | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-raz-adls               | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-raz-chained-plugins    | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-raz-hook-abfs          | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-raz-intg               | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-raz-processor          | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-rms-common             | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-rms-hive               | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-rms-plugins-common     | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-rms-webapp             | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-sampleapp-plugin       | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-schema-registry-plugin | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-solr-plugin            | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-solr-plugin-shim       | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin           | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin-shim      | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-storm-plugin           | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-storm-plugin-shim      | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-tagsync                | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-tools                  | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-util                   | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-yarn-plugin            | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | ranger-yarn-plugin-shim       | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | sample-client                 | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | sampleapp                     | 2.4.0.7.1.9.2-10 |
| Ranger  | org.apache.ranger | shaded-raz-hook-abfs          | 2.4.0.7.1.9.2-10 |

| Project | groupId           | artifactId                               | version           |
|---------|-------------------|------------------------------------------|-------------------|
| Ranger  | org.apache.ranger | ugsync-util                              | 2.4.0.7.1.9.2-10  |
| Ranger  | org.apache.ranger | unixauthclient                           | 2.4.0.7.1.9.2-10  |
| Ranger  | org.apache.ranger | unixauthservice                          | 2.4.0.7.1.9.2-10  |
| Ranger  | org.apache.ranger | unixusersync                             | 2.4.0.7.1.9.2-10  |
| Solr    | org.apache.solr   | solr-analysis-extras                     | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-analytics                           | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-cell                                | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-core                                | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-dataimporthandler                   | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-dataimporthandler-extras            | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-gcs-repository                      | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-jaegertracer-configurator           | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-langid                              | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-ltr                                 | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-prometheus-exporter                 | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-s3-repository                       | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-security-util                       | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-solrj                               | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-test-framework                      | 8.11.2.7.1.9.2-10 |
| Solr    | org.apache.solr   | solr-velocity                            | 8.11.2.7.1.9.2-10 |
| Spark   | org.apache.spark  | spark-avro_2.11                          | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-catalyst_2.11                      | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-core_2.11                          | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-graphx_2.11                        | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-hadoop-cloud_2.11                  | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-hive_2.11                          | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-kubernetes_2.11                    | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-kvstore_2.11                       | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-launcher_2.11                      | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-mllib-local_2.11                   | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-mllib_2.11                         | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-network-common_2.11                | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-network-shuffle_2.11               | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-network-yarn_2.11                  | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-repl_2.11                          | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-sketch_2.11                        | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-sql-kafka-0-10_2.11                | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-sql_2.11                           | 2.4.8.7.1.9.2-10  |
| Spark   | org.apache.spark  | spark-streaming-kafka-0-10-assembly_2.11 | 2.4.8.7.1.9.2-10  |

| Project   | groupId              | artifactId                           | version          |
|-----------|----------------------|--------------------------------------|------------------|
| Spark     | org.apache.spark     | spark-streaming-kafka-0-10_2.11      | 2.4.8.7.1.9.2-10 |
| Spark     | org.apache.spark     | spark-streaming_2.11                 | 2.4.8.7.1.9.2-10 |
| Spark     | org.apache.spark     | spark-tags_2.11                      | 2.4.8.7.1.9.2-10 |
| Spark     | org.apache.spark     | spark-token-provider-kafka-0-10_2.11 | 2.4.8.7.1.9.2-10 |
| Spark     | org.apache.spark     | spark-unsafe_2.11                    | 2.4.8.7.1.9.2-10 |
| Spark     | org.apache.spark     | spark-yarn_2.11                      | 2.4.8.7.1.9.2-10 |
| Sqoop     | org.apache.sqoop     | sqoop                                | 1.4.7.7.1.9.2-10 |
| Sqoop     | org.apache.sqoop     | sqoop-test                           | 1.4.7.7.1.9.2-10 |
| Tez       | org.apache.tez       | hadoop-shim                          | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | hadoop-shim-2.8                      | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-api                              | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-aux-services                     | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-common                           | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-dag                              | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-examples                         | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-ext-service-tests                | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-history-parser                   | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-javadoc-tools                    | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-job-analyzer                     | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-mapreduce                        | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-protobuf-history-plugin          | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-runtime-internals                | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-runtime-library                  | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-tests                            | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-yarn-timeline-cache-plugin       | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history            | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history-with-acls  | 0.9.1.7.1.9.2-10 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history-with-fs    | 0.9.1.7.1.9.2-10 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-angular                     | 0.8.2.7.1.9.2-10 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-display                     | 0.8.2.7.1.9.2-10 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-interpreter                 | 0.8.2.7.1.9.2-10 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-jdbc                        | 0.8.2.7.1.9.2-10 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-jupyter                     | 0.8.2.7.1.9.2-10 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-livy                        | 0.8.2.7.1.9.2-10 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-markdown                    | 0.8.2.7.1.9.2-10 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-server                      | 0.8.2.7.1.9.2-10 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-shell                       | 0.8.2.7.1.9.2-10 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-zengine                     | 0.8.2.7.1.9.2-10 |
| ZooKeeper | org.apache.zookeeper | zookeeper                            | 3.8.1.7.1.9.2-10 |

| Project   | groupId              | artifactId                     | version          |
|-----------|----------------------|--------------------------------|------------------|
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-fatjar       | 3.8.1.7.1.9.2-10 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-loggraph     | 3.8.1.7.1.9.2-10 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-rest         | 3.8.1.7.1.9.2-10 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-zooinspector | 3.8.1.7.1.9.2-10 |
| ZooKeeper | org.apache.zookeeper | zookeeper-it                   | 3.8.1.7.1.9.2-10 |
| ZooKeeper | org.apache.zookeeper | zookeeper-jute                 | 3.8.1.7.1.9.2-10 |
| ZooKeeper | org.apache.zookeeper | zookeeper-prometheus-metrics   | 3.8.1.7.1.9.2-10 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-election     | 3.8.1.7.1.9.2-10 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-lock         | 3.8.1.7.1.9.2-10 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-queue        | 3.8.1.7.1.9.2-10 |

### Runtime 7.1.9.3-4

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          | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-aws-s3-bridge          | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-azure-adls-bridge      | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-classification-updater | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-client-common          | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-client-v1              | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-client-v2              | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-common                 | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-distro                 | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-docs                   | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-graphdb-api            | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-graphdb-common         | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-graphdb-janus          | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-hdfs-bridge            | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-index-repair-tool      | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-intg                   | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-janusgraph-hbase2      | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-notification           | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-plugin-classloader     | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-repository             | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-server-api             | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | atlas-testtools              | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | hbase-bridge                 | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | hbase-bridge-shim            | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | hbase-testing-util           | 3.0.0.7.1.9.3-4 |
| Atlas   | org.apache.atlas | hdfs-model                   | 3.0.0.7.1.9.3-4 |

| Project | groupId            | artifactId                   | version          |
|---------|--------------------|------------------------------|------------------|
| Atlas   | org.apache.atlas   | hive-bridge                  | 3.0.0.7.1.9.3-4  |
| Atlas   | org.apache.atlas   | hive-bridge-shim             | 3.0.0.7.1.9.3-4  |
| Atlas   | org.apache.atlas   | impala-bridge                | 3.0.0.7.1.9.3-4  |
| Atlas   | org.apache.atlas   | impala-bridge-shim           | 3.0.0.7.1.9.3-4  |
| Atlas   | org.apache.atlas   | impala-hook-api              | 3.0.0.7.1.9.3-4  |
| Atlas   | org.apache.atlas   | kafka-bridge                 | 3.0.0.7.1.9.3-4  |
| Atlas   | org.apache.atlas   | kafka-bridge-shim            | 3.0.0.7.1.9.3-4  |
| Atlas   | org.apache.atlas   | navigator-to-atlas           | 3.0.0.7.1.9.3-4  |
| Atlas   | org.apache.atlas   | sample-app                   | 3.0.0.7.1.9.3-4  |
| Atlas   | org.apache.atlas   | sqoop-bridge                 | 3.0.0.7.1.9.3-4  |
| Atlas   | org.apache.atlas   | sqoop-bridge-shim            | 3.0.0.7.1.9.3-4  |
| Avro    | org.apache.avro    | avro                         | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-android                 | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-codegen-test            | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-compiler                | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-grpc                    | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-ipc                     | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-ipc-jetty               | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-ipc-netty               | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-mapred                  | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-maven-plugin            | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-perf                    | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-protobuf                | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-service-archetype       | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-test-custom-conversions | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-thrift                  | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | avro-tools                   | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | trevni-avro                  | 1.11.1.7.1.9.3-4 |
| Avro    | org.apache.avro    | trevni-core                  | 1.11.1.7.1.9.3-4 |
| Calcite | org.apache.calcite | calcite-babel                | 1.19.0.7.1.9.3-4 |
| Calcite | org.apache.calcite | calcite-core                 | 1.19.0.7.1.9.3-4 |
| Calcite | org.apache.calcite | calcite-druid                | 1.19.0.7.1.9.3-4 |
| Calcite | org.apache.calcite | calcite-linq4j               | 1.19.0.7.1.9.3-4 |
| Calcite | org.apache.calcite | calcite-server               | 1.19.0.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop  | hadoop-aliyun                | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop  | hadoop-annotations           | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop  | hadoop-archive-logs          | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop  | hadoop-archives              | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop  | hadoop-assemblies            | 3.1.1.7.1.9.3-4  |



| Project | groupId           | artifactId                         | version         |
|---------|-------------------|------------------------------------|-----------------|
| Hadoop  | org.apache.hadoop | hadoop-auth                        | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-aws                         | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-azure                       | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-azure-datalake              | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-build-tools                 | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-client                      | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-client-api                  | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-client-integration-tests    | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-client-minicluster          | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-client-runtime              | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-cloud-storage               | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-common                      | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-datajoin                    | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-distcp                      | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-extras                      | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-fs2img                      | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-gridmix                     | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs                        | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-client                 | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-httpfs                 | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-native-client          | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-nfs                    | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-rbf                    | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-kafka                       | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-kms                         | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-app        | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-common     | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-core       | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-hs         | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-hs-plugins | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-jobclient  | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-nativetask | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-shuffle    | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-uploader   | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-examples          | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-maven-plugins               | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-minicluster                 | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-minikdc                     | 3.1.1.7.1.9.3-4 |
| Hadoop  | org.apache.hadoop | hadoop-nfs                         | 3.1.1.7.1.9.3-4 |

| Project | groupId           | artifactId                                        | version          |
|---------|-------------------|---------------------------------------------------|------------------|
| Hadoop  | org.apache.hadoop | hadoop-openstack                                  | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-resourceestimator                          | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-rumen                                      | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-sls                                        | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-streaming                                  | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-tools-dist                                 | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-api                                   | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-distributedshell         | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-unmanaged-am-launcher    | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-client                                | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-common                                | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-registry                              | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-applicationhistoryservice      | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-common                         | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-nodemanager                    | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-resourcemanager                | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-router                         | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-sharedcachemanager             | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-tests                          | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timeline-pluginstorage         | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice                | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-client   | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-common   | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-server-2 | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-tests    | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-web-proxy                      | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-api                          | 3.1.1.7.1.9.3-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-core                         | 3.1.1.7.1.9.3-4  |
| HBase   | org.apache.hbase  | hbase-annotations                                 | 2.4.17.7.1.9.3-4 |
| HBase   | org.apache.hbase  | hbase-asyncfs                                     | 2.4.17.7.1.9.3-4 |
| HBase   | org.apache.hbase  | hbase-checkstyle                                  | 2.4.17.7.1.9.3-4 |
| HBase   | org.apache.hbase  | hbase-client                                      | 2.4.17.7.1.9.3-4 |
| HBase   | org.apache.hbase  | hbase-client-project                              | 2.4.17.7.1.9.3-4 |
| HBase   | org.apache.hbase  | hbase-common                                      | 2.4.17.7.1.9.3-4 |
| HBase   | org.apache.hbase  | hbase-endpoint                                    | 2.4.17.7.1.9.3-4 |
| HBase   | org.apache.hbase  | hbase-examples                                    | 2.4.17.7.1.9.3-4 |
| HBase   | org.apache.hbase  | hbase-external-blockcache                         | 2.4.17.7.1.9.3-4 |
| HBase   | org.apache.hbase  | hbase-hadoop-compat                               | 2.4.17.7.1.9.3-4 |
| HBase   | org.apache.hbase  | hbase-hadoop2-compat                              | 2.4.17.7.1.9.3-4 |

| Project | groupId          | artifactId                       | version            |
|---------|------------------|----------------------------------|--------------------|
| HBase   | org.apache.hbase | hbase-hbtop                      | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-http                       | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-it                         | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-logging                    | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-mapreduce                  | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-metrics                    | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-metrics-api                | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-procedure                  | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-protocol                   | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-protocol-shaded            | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-replication                | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-resource-bundle            | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-rest                       | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-rsgroup                    | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-server                     | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-shaded-client              | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-shaded-client-byo-hadoop   | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-shaded-client-project      | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-shaded-mapreduce           | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-shaded-testing-util        | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-shaded-testing-util-tester | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-shell                      | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-testing-util               | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-thrift                     | 2.4.17.7.1.9.3-4   |
| HBase   | org.apache.hbase | hbase-zookeeper                  | 2.4.17.7.1.9.3-4   |
| Hive    | org.apache.hive  | hive-beeline                     | 3.1.3000.7.1.9.3-4 |
| Hive    | org.apache.hive  | hive-blobstore                   | 3.1.3000.7.1.9.3-4 |
| Hive    | org.apache.hive  | hive-classification              | 3.1.3000.7.1.9.3-4 |
| Hive    | org.apache.hive  | hive-cli                         | 3.1.3000.7.1.9.3-4 |
| Hive    | org.apache.hive  | hive-common                      | 3.1.3000.7.1.9.3-4 |
| Hive    | org.apache.hive  | hive-contrib                     | 3.1.3000.7.1.9.3-4 |
| Hive    | org.apache.hive  | hive-exec                        | 3.1.3000.7.1.9.3-4 |
| Hive    | org.apache.hive  | hive-hbase-handler               | 3.1.3000.7.1.9.3-4 |
| Hive    | org.apache.hive  | hive-hcatalog-it-unit            | 3.1.3000.7.1.9.3-4 |
| Hive    | org.apache.hive  | hive-hplsql                      | 3.1.3000.7.1.9.3-4 |
| Hive    | org.apache.hive  | hive-it-custom-serde             | 3.1.3000.7.1.9.3-4 |
| Hive    | org.apache.hive  | hive-it-minikdc                  | 3.1.3000.7.1.9.3-4 |
| Hive    | org.apache.hive  | hive-it-qfile                    | 3.1.3000.7.1.9.3-4 |
| Hive    | org.apache.hive  | hive-it-qfile-kudu               | 3.1.3000.7.1.9.3-4 |

| Project                  | groupId              | artifactId                               | version            |
|--------------------------|----------------------|------------------------------------------|--------------------|
| Hive                     | org.apache.hive      | hive-it-test-serde                       | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-it-unit                             | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-it-unit-hadoop2                     | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-it-util                             | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-jdbc                                | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-jdbc-handler                        | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-jmh                                 | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-kryo-registrator                    | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-kudu-handler                        | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-llap-client                         | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-llap-common                         | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-llap-ext-client                     | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-llap-server                         | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-llap-tez                            | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-metastore                           | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-pre-upgrade                         | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-serde                               | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-service                             | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-service-rpc                         | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-shims                               | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-spark-client                        | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-standalone-metastore                | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-storage-api                         | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-streaming                           | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-testutils                           | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | hive-vector-code-gen                     | 3.1.3000.7.1.9.3-4 |
| Hive                     | org.apache.hive      | kafka-handler                            | 3.1.3000.7.1.9.3-4 |
| Hive Warehouse Connector | com.hortonworks.hive | hive-warehouse-connector_2.11            | 1.0.0.7.1.9.3-4    |
| Kafka                    | org.apache.kafka     | ci                                       | 3.4.1.7.1.9.3-4    |
| Kafka                    | org.apache.kafka     | connect                                  | 3.4.1.7.1.9.3-4    |
| Kafka                    | org.apache.kafka     | connect-api                              | 3.4.1.7.1.9.3-4    |
| Kafka                    | org.apache.kafka     | connect-basic-auth-extension             | 3.4.1.7.1.9.3-4    |
| Kafka                    | org.apache.kafka     | connect-cloudera-authorization-extension | 3.4.1.7.1.9.3-4    |
| Kafka                    | org.apache.kafka     | connect-cloudera-common                  | 3.4.1.7.1.9.3-4    |
| Kafka                    | org.apache.kafka     | connect-cloudera-secret-storage          | 3.4.1.7.1.9.3-4    |
| Kafka                    | org.apache.kafka     | connect-cloudera-security-policies       | 3.4.1.7.1.9.3-4    |
| Kafka                    | org.apache.kafka     | connect-file                             | 3.4.1.7.1.9.3-4    |
| Kafka                    | org.apache.kafka     | connect-json                             | 3.4.1.7.1.9.3-4    |
| Kafka                    | org.apache.kafka     | connect-mirror                           | 3.4.1.7.1.9.3-4    |

| Project | groupId          | artifactId                              | version         |
|---------|------------------|-----------------------------------------|-----------------|
| Kafka   | org.apache.kafka | connect-mirror-client                   | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | connect-runtime                         | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | connect-transforms                      | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | generator                               | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-clients                           | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-cloudera-metrics-reporter_2.12    | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-cloudera-metrics-reporter_2.13    | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-cloudera-plugins                  | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-examples                          | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-group-coordinator                 | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-log4j-appender                    | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-metadata                          | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-raft                              | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-server-common                     | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-shell                             | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-storage                           | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-storage-api                       | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams                           | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-examples                  | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-scala_2.12                | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-scala_2.13                | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-test-utils                | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0100 | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0101 | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0102 | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0110 | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-10   | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-11   | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-20   | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-21   | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-22   | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-23   | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-24   | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-25   | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-26   | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-27   | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-28   | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-30   | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-31   | 3.4.1.7.1.9.3-4 |

| Project | groupId          | artifactId                                        | version         |
|---------|------------------|---------------------------------------------------|-----------------|
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-32             | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-33             | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka-tools                                       | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka_2.12                                        | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | kafka_2.13                                        | 3.4.1.7.1.9.3-4 |
| Kafka   | org.apache.kafka | trogdor                                           | 3.4.1.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-adapter                                   | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-admin-ui                                  | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-applications                              | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-cloud-bindings                            | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-demo-ldap                                 | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-demo-ldap-launcher                        | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-discovery-ambari                          | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-discovery-cm                              | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-docker                                    | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-i18n                                      | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-i18n-logging-log4j                        | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-i18n-logging-sl4j                         | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-performance-test                          | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-ha                               | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-common        | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-concat        | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-hadoop-groups | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-pseudo        | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-regex         | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-switchcase    | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-jersey                           | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-rewrite                          | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-common                   | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-func-hostmap-static      | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-func-inbound-query-param | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-func-service-registry    | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-step-encrypt-uri         | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-step-secure-query        | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-security-authc-anon              | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-security-authz-acls              | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-security-authz-composite         | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-security-clientcert              | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox  | gateway-provider-security-hadoopauth              | 1.3.0.7.1.9.3-4 |

| Project | groupId         | artifactId                          | version         |
|---------|-----------------|-------------------------------------|-----------------|
| Knox    | org.apache.knox | gateway-provider-security-jwt       | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-provider-security-pac4j     | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-provider-security-preauth   | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-provider-security-shiro     | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-provider-security-webappsec | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-release                     | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-server                      | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-server-launcher             | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-server-xforwarded-filter    | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-admin               | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-as                  | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-definitions         | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-hashicorp-vault     | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-hbase               | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-health              | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-hive                | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-idbroker            | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-impala              | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-jkg                 | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-knoxsso             | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-knoxssout           | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-knoxtoken           | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-livy                | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-metadata            | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-nifi                | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-nifi-registry       | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-remoteconfig        | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-rm                  | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-session             | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-storm               | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-test                | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-tgs                 | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-vault               | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-service-webhdfs             | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-shell                       | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-shell-launcher              | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-shell-release               | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-shell-samples               | 1.3.0.7.1.9.3-4 |
| Knox    | org.apache.knox | gateway-spi                         | 1.3.0.7.1.9.3-4 |

| Project | groupId         | artifactId                  | version          |
|---------|-----------------|-----------------------------|------------------|
| Knox    | org.apache.knox | gateway-test                | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | gateway-test-idbroker       | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | gateway-test-release-utils  | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | gateway-test-utils          | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | gateway-topology-hadoop-xml | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | gateway-topology-simple     | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | gateway-util-common         | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | gateway-util-configinjector | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | gateway-util-launcher       | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | gateway-util-urltemplate    | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | hadoop-examples             | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | knox-cli-launcher           | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | knox-homepage-ui            | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | knox-token-generation-ui    | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | knox-token-management-ui    | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | webhdfs-kerb-test           | 1.3.0.7.1.9.3-4  |
| Knox    | org.apache.knox | webhdfs-test                | 1.3.0.7.1.9.3-4  |
| Kudu    | org.apache.kudu | kudu-backup-tools           | 1.17.0.7.1.9.3-4 |
| Kudu    | org.apache.kudu | kudu-backup2_2.11           | 1.17.0.7.1.9.3-4 |
| Kudu    | org.apache.kudu | kudu-backup3_2.12           | 1.17.0.7.1.9.3-4 |
| Kudu    | org.apache.kudu | kudu-client                 | 1.17.0.7.1.9.3-4 |
| Kudu    | org.apache.kudu | kudu-hive                   | 1.17.0.7.1.9.3-4 |
| Kudu    | org.apache.kudu | kudu-spark2-tools_2.11      | 1.17.0.7.1.9.3-4 |
| Kudu    | org.apache.kudu | kudu-spark2_2.11            | 1.17.0.7.1.9.3-4 |
| Kudu    | org.apache.kudu | kudu-spark3-tools_2.12      | 1.17.0.7.1.9.3-4 |
| Kudu    | org.apache.kudu | kudu-spark3_2.12            | 1.17.0.7.1.9.3-4 |
| Kudu    | org.apache.kudu | kudu-test-utils             | 1.17.0.7.1.9.3-4 |
| Livy    | org.apache.livy | livy-api                    | 0.7.2.7.1.9.3-4  |
| Livy    | org.apache.livy | livy-client-common          | 0.7.2.7.1.9.3-4  |
| Livy    | org.apache.livy | livy-client-http            | 0.7.2.7.1.9.3-4  |
| Livy    | org.apache.livy | livy-core_2.11              | 0.7.2.7.1.9.3-4  |
| Livy    | org.apache.livy | livy-examples               | 0.7.2.7.1.9.3-4  |
| Livy    | org.apache.livy | livy-integration-test       | 0.7.2.7.1.9.3-4  |
| Livy    | org.apache.livy | livy-repl_2.11              | 0.7.2.7.1.9.3-4  |
| Livy    | org.apache.livy | livy-rsc                    | 0.7.2.7.1.9.3-4  |
| Livy    | org.apache.livy | livy-scala-api_2.11         | 0.7.2.7.1.9.3-4  |
| Livy    | org.apache.livy | livy-server                 | 0.7.2.7.1.9.3-4  |
| Livy    | org.apache.livy | livy-test-lib               | 0.7.2.7.1.9.3-4  |
| Livy    | org.apache.livy | livy-thriftserver           | 0.7.2.7.1.9.3-4  |



| Project | groupId           | artifactId                  | version          |
|---------|-------------------|-----------------------------|------------------|
| Livy    | org.apache.livy   | livy-thriftserver-session   | 0.7.2.7.1.9.3-4  |
| Lucene  | org.apache.lucene | lucene-analyzers-common     | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-icu        | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-kuromoji   | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-morfologik | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-nori       | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-openslp    | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-phonetic   | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-smartcn    | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-stempel    | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-backward-codecs      | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-benchmark            | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-classification       | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-codecs               | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-core                 | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-demo                 | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-expressions          | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-facet                | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-grouping             | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-highlighter          | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-join                 | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-memory               | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-misc                 | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-monitor              | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-queries              | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-queryparser          | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-replicator           | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-sandbox              | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-spatial-extras       | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-spatial3d            | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-suggest              | 8.11.2.7.1.9.3-4 |
| Lucene  | org.apache.lucene | lucene-test-framework       | 8.11.2.7.1.9.3-4 |
| Oozie   | org.apache.oozie  | oozie-client                | 5.1.0.7.1.9.3-4  |
| Oozie   | org.apache.oozie  | oozie-core                  | 5.1.0.7.1.9.3-4  |
| Oozie   | org.apache.oozie  | oozie-distro                | 5.1.0.7.1.9.3-4  |
| Oozie   | org.apache.oozie  | oozie-examples              | 5.1.0.7.1.9.3-4  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-api        | 5.1.0.7.1.9.3-4  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-client     | 5.1.0.7.1.9.3-4  |
| Oozie   | org.apache.oozie  | oozie-server                | 5.1.0.7.1.9.3-4  |

| Project | groupId          | artifactId                     | version         |
|---------|------------------|--------------------------------|-----------------|
| Oozie   | org.apache.oozie | oozie-sharelib-distcp          | 5.1.0.7.1.9.3-4 |
| Oozie   | org.apache.oozie | oozie-sharelib-git             | 5.1.0.7.1.9.3-4 |
| Oozie   | org.apache.oozie | oozie-sharelib-hcatalog        | 5.1.0.7.1.9.3-4 |
| Oozie   | org.apache.oozie | oozie-sharelib-hive            | 5.1.0.7.1.9.3-4 |
| Oozie   | org.apache.oozie | oozie-sharelib-hive2           | 5.1.0.7.1.9.3-4 |
| Oozie   | org.apache.oozie | oozie-sharelib-oozie           | 5.1.0.7.1.9.3-4 |
| Oozie   | org.apache.oozie | oozie-sharelib-spark           | 5.1.0.7.1.9.3-4 |
| Oozie   | org.apache.oozie | oozie-sharelib-spark3          | 5.1.0.7.1.9.3-4 |
| Oozie   | org.apache.oozie | oozie-sharelib-sqoop           | 5.1.0.7.1.9.3-4 |
| Oozie   | org.apache.oozie | oozie-sharelib-streaming       | 5.1.0.7.1.9.3-4 |
| Oozie   | org.apache.oozie | oozie-tools                    | 5.1.0.7.1.9.3-4 |
| Oozie   | org.apache.oozie | oozie-zookeeper-security-tests | 5.1.0.7.1.9.3-4 |
| ORC     | org.apache.orc   | orc-core                       | 1.5.1.7.1.9.3-4 |
| ORC     | org.apache.orc   | orc-examples                   | 1.5.1.7.1.9.3-4 |
| ORC     | org.apache.orc   | orc-mapreduce                  | 1.5.1.7.1.9.3-4 |
| ORC     | org.apache.orc   | orc-shims                      | 1.5.1.7.1.9.3-4 |
| ORC     | org.apache.orc   | orc-tools                      | 1.5.1.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-annotation-processing     | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-client                    | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-common                    | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-config                    | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-container-service         | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-docs                      | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-erasurecode               | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-client  | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-server  | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-test    | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-interface-admin           | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-interface-client          | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-interface-server          | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-managed-rocksdbs          | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-rocks-native              | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-server-framework          | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-server-scm                | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-test-utils                | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | hdds-tools                     | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | mini-chaos-tests               | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | ozone-client                   | 1.2.0.7.1.9.3-4 |
| Ozone   | org.apache.ozone | ozone-common                   | 1.2.0.7.1.9.3-4 |

| Project | groupId            | artifactId                        | version           |
|---------|--------------------|-----------------------------------|-------------------|
| Ozone   | org.apache.ozone   | ozone-csi                         | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-datanode                    | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-dist                        | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-filessystem                 | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-filessystem-common          | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-filessystem-hadoop2         | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-filessystem-hadoop3         | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-filessystem-shaded          | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-httpsgateway                | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-insight                     | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-integration-test            | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-interface-client            | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-interface-storage           | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-manager                     | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-network-tests               | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-recon                       | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-reconcodegen                | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-s3-secret-store             | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-s3gateway                   | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | ozone-tools                       | 1.2.0.7.1.9.3-4   |
| Ozone   | org.apache.ozone   | rocksdb-checkpoint-differ         | 1.2.0.7.1.9.3-4   |
| Parquet | org.apache.parquet | parquet-avro                      | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-cascading                 | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-cascading3                | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-column                    | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-common                    | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-encoding                  | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-format-structures         | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-generator                 | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-hadoop                    | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-hadoop-bundle             | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-jackson                   | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-pig                       | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-pig-bundle                | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-protobuf                  | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-scala_2.10                | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-thrift                    | 1.10.99.7.1.9.3-4 |
| Parquet | org.apache.parquet | parquet-tools                     | 1.10.99.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-client-embedded-hbase-2.4 | 5.1.1.7.1.9.3-4   |

| Project | groupId            | artifactId                             | version         |
|---------|--------------------|----------------------------------------|-----------------|
| Phoenix | org.apache.phoenix | phoenix-client-hbase-2.4               | 5.1.1.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-connectors-phoenix5-compatible | 6.0.0.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-core                           | 5.1.1.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.1.6         | 5.1.1.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.2.5         | 5.1.1.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.3.0         | 5.1.1.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.4.0         | 5.1.1.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.4.1         | 5.1.1.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-pherf                          | 5.1.1.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-queryserver                    | 6.0.0.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-client             | 6.0.0.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-it                 | 6.0.0.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-load-balancer      | 6.0.0.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-orchestrator       | 6.0.0.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-server-hbase-2.4               | 5.1.1.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix-tracing-webapp                 | 5.1.1.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix5-hive                          | 6.0.0.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix5-hive-shaded                   | 6.0.0.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix5-spark                         | 6.0.0.7.1.9.3-4 |
| Phoenix | org.apache.phoenix | phoenix5-spark-shaded                  | 6.0.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | conditions-enrichers                   | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | credentialbuilder                      | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | embeddedwebserver                      | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | jisql                                  | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ldapconfigcheck                        | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ranger-adls-plugin                     | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ranger-atlas-plugin                    | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ranger-atlas-plugin-shim               | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ranger-authn                           | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ranger-common-ha                       | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ranger-distro                          | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ranger-examples-distro                 | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ranger-hbase-plugin                    | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ranger-hbase-plugin-shim               | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ranger-hdfs-plugin                     | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ranger-hdfs-plugin-shim                | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ranger-hive-plugin                     | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ranger-hive-plugin-shim                | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger  | ranger-intg                            | 2.4.0.7.1.9.3-4 |

| Project | groupId           | artifactId                    | version         |
|---------|-------------------|-------------------------------|-----------------|
| Ranger  | org.apache.ranger | ranger-kafka-connect-plugin   | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-kafka-plugin           | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-kafka-plugin-shim      | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-kms                    | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-kms-plugin             | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-kms-plugin-shim        | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-knox-plugin            | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-knox-plugin-shim       | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-kudu-plugin            | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-kylin-plugin           | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-kylin-plugin-shim      | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-metrics                | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-nifi-plugin            | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-nifi-registry-plugin   | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-ozone-plugin           | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-ozone-plugin-shim      | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-plugin-classloader     | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-plugins-audit          | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-plugins-common         | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-plugins-cred           | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-plugins-installer      | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-policymigration        | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-raz-adls               | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-raz-chained-plugins    | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-raz-hook-abfs          | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-raz-intg               | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-raz-processor          | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-rms-common             | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-rms-hive               | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-rms-plugins-common     | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-rms-webapp             | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-sampleapp-plugin       | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-schema-registry-plugin | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-solr-plugin            | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-solr-plugin-shim       | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin           | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin-shim      | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-storm-plugin           | 2.4.0.7.1.9.3-4 |
| Ranger  | org.apache.ranger | ranger-storm-plugin-shim      | 2.4.0.7.1.9.3-4 |

| Project | groupId           | artifactId                     | version          |
|---------|-------------------|--------------------------------|------------------|
| Ranger  | org.apache.ranger | ranger-tagsync                 | 2.4.0.7.1.9.3-4  |
| Ranger  | org.apache.ranger | ranger-tools                   | 2.4.0.7.1.9.3-4  |
| Ranger  | org.apache.ranger | ranger-util                    | 2.4.0.7.1.9.3-4  |
| Ranger  | org.apache.ranger | ranger-yarn-plugin             | 2.4.0.7.1.9.3-4  |
| Ranger  | org.apache.ranger | ranger-yarn-plugin-shim        | 2.4.0.7.1.9.3-4  |
| Ranger  | org.apache.ranger | sample-client                  | 2.4.0.7.1.9.3-4  |
| Ranger  | org.apache.ranger | sampleapp                      | 2.4.0.7.1.9.3-4  |
| Ranger  | org.apache.ranger | shaded-raz-hook-abfs           | 2.4.0.7.1.9.3-4  |
| Ranger  | org.apache.ranger | ugsync-util                    | 2.4.0.7.1.9.3-4  |
| Ranger  | org.apache.ranger | unixauthclient                 | 2.4.0.7.1.9.3-4  |
| Ranger  | org.apache.ranger | unixauthservice                | 2.4.0.7.1.9.3-4  |
| Ranger  | org.apache.ranger | unixusersync                   | 2.4.0.7.1.9.3-4  |
| Solr    | org.apache.solr   | solr-analysis-extras           | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-analytics                 | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-cell                      | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-core                      | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-dataimporthandler         | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-dataimporthandler-extras  | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-gcs-repository            | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-jaegertracer-configurator | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-langid                    | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-ltr                       | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-prometheus-exporter       | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-s3-repository             | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-security-util             | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-solrj                     | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-test-framework            | 8.11.2.7.1.9.3-4 |
| Solr    | org.apache.solr   | solr-velocity                  | 8.11.2.7.1.9.3-4 |
| Spark   | org.apache.spark  | spark-avro_2.11                | 2.4.8.7.1.9.3-4  |
| Spark   | org.apache.spark  | spark-catalyst_2.11            | 2.4.8.7.1.9.3-4  |
| Spark   | org.apache.spark  | spark-core_2.11                | 2.4.8.7.1.9.3-4  |
| Spark   | org.apache.spark  | spark-graphx_2.11              | 2.4.8.7.1.9.3-4  |
| Spark   | org.apache.spark  | spark-hadoop-cloud_2.11        | 2.4.8.7.1.9.3-4  |
| Spark   | org.apache.spark  | spark-hive_2.11                | 2.4.8.7.1.9.3-4  |
| Spark   | org.apache.spark  | spark-kubernetes_2.11          | 2.4.8.7.1.9.3-4  |
| Spark   | org.apache.spark  | spark-kvstore_2.11             | 2.4.8.7.1.9.3-4  |
| Spark   | org.apache.spark  | spark-launcher_2.11            | 2.4.8.7.1.9.3-4  |
| Spark   | org.apache.spark  | spark-mllib-local_2.11         | 2.4.8.7.1.9.3-4  |
| Spark   | org.apache.spark  | spark-mllib_2.11               | 2.4.8.7.1.9.3-4  |

| Project  | groupId             | artifactId                               | version         |
|----------|---------------------|------------------------------------------|-----------------|
| Spark    | org.apache.spark    | spark-network-common_2.11                | 2.4.8.7.1.9.3-4 |
| Spark    | org.apache.spark    | spark-network-shuffle_2.11               | 2.4.8.7.1.9.3-4 |
| Spark    | org.apache.spark    | spark-network-yarn_2.11                  | 2.4.8.7.1.9.3-4 |
| Spark    | org.apache.spark    | spark-repl_2.11                          | 2.4.8.7.1.9.3-4 |
| Spark    | org.apache.spark    | spark-sketch_2.11                        | 2.4.8.7.1.9.3-4 |
| Spark    | org.apache.spark    | spark-sql-kafka-0-10_2.11                | 2.4.8.7.1.9.3-4 |
| Spark    | org.apache.spark    | spark-sql_2.11                           | 2.4.8.7.1.9.3-4 |
| Spark    | org.apache.spark    | spark-streaming-kafka-0-10-assembly_2.11 | 2.4.8.7.1.9.3-4 |
| Spark    | org.apache.spark    | spark-streaming-kafka-0-10_2.11          | 2.4.8.7.1.9.3-4 |
| Spark    | org.apache.spark    | spark-streaming_2.11                     | 2.4.8.7.1.9.3-4 |
| Spark    | org.apache.spark    | spark-tags_2.11                          | 2.4.8.7.1.9.3-4 |
| Spark    | org.apache.spark    | spark-token-provider-kafka-0-10_2.11     | 2.4.8.7.1.9.3-4 |
| Spark    | org.apache.spark    | spark-unsafe_2.11                        | 2.4.8.7.1.9.3-4 |
| Spark    | org.apache.spark    | spark-yarn_2.11                          | 2.4.8.7.1.9.3-4 |
| Sqoop    | org.apache.sqoop    | sqoop                                    | 1.4.7.7.1.9.3-4 |
| Sqoop    | org.apache.sqoop    | sqoop-test                               | 1.4.7.7.1.9.3-4 |
| Tez      | org.apache.tez      | hadoop-shim                              | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | hadoop-shim-2.8                          | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-api                                  | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-aux-services                         | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-common                               | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-dag                                  | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-examples                             | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-ext-service-tests                    | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-history-parser                       | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-javadoc-tools                        | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-job-analyzer                         | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-mapreduce                            | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-protobuf-history-plugin              | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-runtime-internals                    | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-runtime-library                      | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-tests                                | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-yarn-timeline-cache-plugin           | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-yarn-timeline-history                | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-yarn-timeline-history-with-acls      | 0.9.1.7.1.9.3-4 |
| Tez      | org.apache.tez      | tez-yarn-timeline-history-with-fs        | 0.9.1.7.1.9.3-4 |
| Zeppelin | org.apache.zeppelin | zeppelin-angular                         | 0.8.2.7.1.9.3-4 |
| Zeppelin | org.apache.zeppelin | zeppelin-display                         | 0.8.2.7.1.9.3-4 |
| Zeppelin | org.apache.zeppelin | zeppelin-interpreter                     | 0.8.2.7.1.9.3-4 |

| Project   | groupId              | artifactId                     | version         |
|-----------|----------------------|--------------------------------|-----------------|
| Zeppelin  | org.apache.zeppelin  | zeppelin-jdbc                  | 0.8.2.7.1.9.3-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-jupyter               | 0.8.2.7.1.9.3-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-livy                  | 0.8.2.7.1.9.3-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-markdown              | 0.8.2.7.1.9.3-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-server                | 0.8.2.7.1.9.3-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-shell                 | 0.8.2.7.1.9.3-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-zengine               | 0.8.2.7.1.9.3-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper                      | 3.8.1.7.1.9.3-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-fatjar       | 3.8.1.7.1.9.3-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-loggraph     | 3.8.1.7.1.9.3-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-rest         | 3.8.1.7.1.9.3-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-zooinspector | 3.8.1.7.1.9.3-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-it                   | 3.8.1.7.1.9.3-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-jute                 | 3.8.1.7.1.9.3-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-prometheus-metrics   | 3.8.1.7.1.9.3-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-election     | 3.8.1.7.1.9.3-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-lock         | 3.8.1.7.1.9.3-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-queue        | 3.8.1.7.1.9.3-4 |

## Runtime 7.1.9.4-4

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          | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-aws-s3-bridge          | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-azure-adls-bridge      | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-classification-updater | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-client-common          | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-client-v1              | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-client-v2              | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-common                 | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-distro                 | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-docs                   | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-graphdb-api            | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-graphdb-common         | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-graphdb-janus          | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-hdfs-bridge            | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-index-repair-tool      | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-intg                   | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-janusgraph-hbase2      | 3.0.0.7.1.9.4-4 |
| Atlas   | org.apache.atlas | atlas-notification           | 3.0.0.7.1.9.4-4 |



| Project | groupId            | artifactId                   | version          |
|---------|--------------------|------------------------------|------------------|
| Atlas   | org.apache.atlas   | atlas-plugin-classloader     | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | atlas-repository             | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | atlas-server-api             | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | atlas-testtools              | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | hbase-bridge                 | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | hbase-bridge-shim            | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | hbase-testing-util           | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | hdfs-model                   | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | hive-bridge                  | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | hive-bridge-shim             | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | impala-bridge                | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | impala-bridge-shim           | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | impala-hook-api              | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | kafka-bridge                 | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | kafka-bridge-shim            | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | navigator-to-atlas           | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | sample-app                   | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | sqoop-bridge                 | 3.0.0.7.1.9.4-4  |
| Atlas   | org.apache.atlas   | sqoop-bridge-shim            | 3.0.0.7.1.9.4-4  |
| Avro    | org.apache.avro    | avro                         | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-android                 | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-codegen-test            | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-compiler                | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-grpc                    | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-ipc                     | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-ipc-jetty               | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-ipc-netty               | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-mapred                  | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-maven-plugin            | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-perf                    | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-protobuf                | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-service-archetype       | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-test-custom-conversions | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-thrift                  | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | avro-tools                   | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | trevni-avro                  | 1.11.1.7.1.9.4-4 |
| Avro    | org.apache.avro    | trevni-core                  | 1.11.1.7.1.9.4-4 |
| Calcite | org.apache.calcite | calcite-babel                | 1.19.0.7.1.9.4-4 |
| Calcite | org.apache.calcite | calcite-core                 | 1.19.0.7.1.9.4-4 |

| Project | groupId            | artifactId                         | version          |
|---------|--------------------|------------------------------------|------------------|
| Calcite | org.apache.calcite | calcite-druid                      | 1.19.0.7.1.9.4-4 |
| Calcite | org.apache.calcite | calcite-linq4j                     | 1.19.0.7.1.9.4-4 |
| Calcite | org.apache.calcite | calcite-server                     | 1.19.0.7.1.9.4-4 |
| Hadoop  | org.apache.hadoop  | hadoop-aliyun                      | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-annotations                 | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-archive-logs                | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-archives                    | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-assemblies                  | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-auth                        | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-aws                         | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-azure                       | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-azure-datalake              | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-build-tools                 | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-client                      | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-client-api                  | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-client-integration-tests    | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-client-minicluster          | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-client-runtime              | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-cloud-storage               | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-common                      | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-datajoin                    | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-distcp                      | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-extras                      | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-fs2img                      | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-gridmix                     | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-hdfs                        | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-hdfs-client                 | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-hdfs-httpfs                 | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-hdfs-native-client          | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-hdfs-nfs                    | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-hdfs-rbf                    | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-kafka                       | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-kms                         | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-mapreduce-client-app        | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-mapreduce-client-common     | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-mapreduce-client-core       | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-mapreduce-client-hs         | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-mapreduce-client-hs-plugins | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop  | hadoop-mapreduce-client-jobclient  | 3.1.1.7.1.9.4-4  |

| Project | groupId           | artifactId                                        | version          |
|---------|-------------------|---------------------------------------------------|------------------|
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-nativetask                | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-shuffle                   | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-uploader                  | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-examples                         | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-maven-plugins                              | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-minicluster                                | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-minikdc                                    | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-nfs                                        | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-openstack                                  | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-resourceestimator                          | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-rumen                                      | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-sls                                        | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-streaming                                  | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-tools-dist                                 | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-api                                   | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-distributedshell         | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-unmanaged-am-launcher    | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-client                                | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-common                                | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-registry                              | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-applicationhistoryservice      | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-common                         | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-nodemanager                    | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-resourcemanager                | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-router                         | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-sharedcachemanager             | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-tests                          | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timeline-pluginstorage         | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice                | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-client   | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-common   | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-server-2 | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-tests    | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-web-proxy                      | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-api                          | 3.1.1.7.1.9.4-4  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-core                         | 3.1.1.7.1.9.4-4  |
| HBase   | org.apache.hbase  | hbase-annotations                                 | 2.4.17.7.1.9.4-4 |
| HBase   | org.apache.hbase  | hbase-asyncfs                                     | 2.4.17.7.1.9.4-4 |
| HBase   | org.apache.hbase  | hbase-checkstyle                                  | 2.4.17.7.1.9.4-4 |

| Project | groupId          | artifactId                       | version            |
|---------|------------------|----------------------------------|--------------------|
| HBase   | org.apache.hbase | hbase-client                     | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-client-project             | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-common                     | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-endpoint                   | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-examples                   | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-external-blockcache        | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-hadoop-compat              | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-hadoop2-compat             | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-hbtop                      | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-http                       | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-it                         | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-logging                    | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-mapreduce                  | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-metrics                    | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-metrics-api                | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-procedure                  | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-protocol                   | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-protocol-shaded            | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-replication                | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-resource-bundle            | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-rest                       | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-rsgroup                    | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-server                     | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-shaded-client              | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-shaded-client-byo-hadoop   | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-shaded-client-project      | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-shaded-mapreduce           | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-shaded-testing-util        | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-shaded-testing-util-tester | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-shell                      | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-testing-util               | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-thrift                     | 2.4.17.7.1.9.4-4   |
| HBase   | org.apache.hbase | hbase-zookeeper                  | 2.4.17.7.1.9.4-4   |
| Hive    | org.apache.hive  | hive-beeline                     | 3.1.3000.7.1.9.4-4 |
| Hive    | org.apache.hive  | hive-blobstore                   | 3.1.3000.7.1.9.4-4 |
| Hive    | org.apache.hive  | hive-classification              | 3.1.3000.7.1.9.4-4 |
| Hive    | org.apache.hive  | hive-cli                         | 3.1.3000.7.1.9.4-4 |
| Hive    | org.apache.hive  | hive-common                      | 3.1.3000.7.1.9.4-4 |
| Hive    | org.apache.hive  | hive-contrib                     | 3.1.3000.7.1.9.4-4 |

| Project                  | groupId              | artifactId                    | version            |
|--------------------------|----------------------|-------------------------------|--------------------|
| Hive                     | org.apache.hive      | hive-exec                     | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-hbase-handler            | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-hcatalog-it-unit         | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-hplsql                   | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-it-custom-serde          | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-it-minikdc               | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-it-qfile                 | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-it-qfile-kudu            | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-it-test-serde            | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-it-unit                  | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-it-unit-hadoop2          | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-it-util                  | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-jdbc                     | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-jdbc-handler             | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-jmh                      | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-kryo-registrator         | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-kudu-handler             | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-llap-client              | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-llap-common              | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-llap-ext-client          | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-llap-server              | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-llap-tez                 | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-metastore                | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-pre-upgrade              | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-serde                    | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-service                  | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-service-rpc              | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-shims                    | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-spark-client             | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-standalone-metastore     | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-storage-api              | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-streaming                | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-testutils                | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | hive-vector-code-gen          | 3.1.3000.7.1.9.4-4 |
| Hive                     | org.apache.hive      | kafka-handler                 | 3.1.3000.7.1.9.4-4 |
| Hive Warehouse Connector | com.hortonworks.hive | hive-warehouse-connector_2.11 | 1.0.0.7.1.9.4-4    |
| Kafka                    | org.apache.kafka     | ci                            | 3.4.1.7.1.9.4-4    |
| Kafka                    | org.apache.kafka     | connect                       | 3.4.1.7.1.9.4-4    |
| Kafka                    | org.apache.kafka     | connect-api                   | 3.4.1.7.1.9.4-4    |

| Project | groupId          | artifactId                               | version         |
|---------|------------------|------------------------------------------|-----------------|
| Kafka   | org.apache.kafka | connect-basic-auth-extension             | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | connect-cloudera-authorization-extension | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | connect-cloudera-common                  | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | connect-cloudera-secret-storage          | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | connect-cloudera-security-policies       | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | connect-file                             | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | connect-json                             | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | connect-mirror                           | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | connect-mirror-client                    | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | connect-runtime                          | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | connect-transforms                       | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | generator                                | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-clients                            | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-cloudera-metrics-reporter_2.12     | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-cloudera-metrics-reporter_2.13     | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-cloudera-plugins                   | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-examples                           | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-group-coordinator                  | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-log4j-appender                     | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-metadata                           | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-raft                               | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-server-common                      | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-shell                              | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-storage                            | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-storage-api                        | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams                            | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-examples                   | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-scala_2.12                 | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-scala_2.13                 | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-test-utils                 | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0100  | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0101  | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0102  | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0110  | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-10    | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-11    | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-20    | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-21    | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-22    | 3.4.1.7.1.9.4-4 |

| Project | groupId          | artifactId                                        | version         |
|---------|------------------|---------------------------------------------------|-----------------|
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-23             | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-24             | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-25             | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-26             | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-27             | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-28             | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-30             | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-31             | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-32             | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-33             | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka-tools                                       | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka_2.12                                        | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | kafka_2.13                                        | 3.4.1.7.1.9.4-4 |
| Kafka   | org.apache.kafka | trogdor                                           | 3.4.1.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-adapter                                   | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-admin-ui                                  | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-applications                              | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-cloud-bindings                            | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-demo-ldap                                 | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-demo-ldap-launcher                        | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-discovery-ambari                          | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-discovery-cm                              | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-docker                                    | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-i18n                                      | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-i18n-logging-log4j                        | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-i18n-logging-sl4j                         | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-performance-test                          | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-provider-ha                               | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-common        | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-concat        | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-hadoop-groups | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-pseudo        | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-regex         | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-switchcase    | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-provider-jersey                           | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-provider-rewrite                          | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-common                   | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-func-hostmap-static      | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-func-inbound-query-param | 1.3.0.7.1.9.4-4 |

| Project | groupId         | artifactId                                     | version         |
|---------|-----------------|------------------------------------------------|-----------------|
| Knox    | org.apache.knox | gateway-provider-rewrite-func-service-registry | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-provider-rewrite-step-encrypt-uri      | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-provider-rewrite-step-secure-query     | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-provider-security-authc-anon           | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-provider-security-authz-acls           | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-provider-security-authz-composite      | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-provider-security-clientcert           | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-provider-security-hadoopauth           | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-provider-security-jwt                  | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-provider-security-pac4j                | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-provider-security-preauth              | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-provider-security-shiro                | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-provider-security-webappsec            | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-release                                | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-server                                 | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-server-launcher                        | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-server-xforwarded-filter               | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-admin                          | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-as                             | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-definitions                    | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-hashicorp-vault                | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-hbase                          | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-health                         | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-hive                           | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-idbroker                       | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-impala                         | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-jkg                            | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-knoxsso                        | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-knoxssout                      | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-knoxtoken                      | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-livy                           | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-metadata                       | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-nifi                           | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-nifi-registry                  | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-remoteconfig                   | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-rm                             | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-session                        | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-storm                          | 1.3.0.7.1.9.4-4 |
| Knox    | org.apache.knox | gateway-service-test                           | 1.3.0.7.1.9.4-4 |



| Project | groupId         | artifactId                  | version          |
|---------|-----------------|-----------------------------|------------------|
| Knox    | org.apache.knox | gateway-service-tgs         | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-service-vault       | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-service-webhdfs     | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-shell               | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-shell-launcher      | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-shell-release       | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-shell-samples       | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-spi                 | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-test                | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-test-idbroker       | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-test-release-utils  | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-test-utils          | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-topology-hadoop-xml | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-topology-simple     | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-util-common         | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-util-configinjector | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-util-launcher       | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | gateway-util-urltemplate    | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | hadoop-examples             | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | knox-cli-launcher           | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | knox-homepage-ui            | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | knox-token-generation-ui    | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | knox-token-management-ui    | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | webhdfs-kerb-test           | 1.3.0.7.1.9.4-4  |
| Knox    | org.apache.knox | webhdfs-test                | 1.3.0.7.1.9.4-4  |
| Kudu    | org.apache.kudu | kudu-backup-tools           | 1.17.0.7.1.9.4-4 |
| Kudu    | org.apache.kudu | kudu-backup2_2.11           | 1.17.0.7.1.9.4-4 |
| Kudu    | org.apache.kudu | kudu-backup3_2.12           | 1.17.0.7.1.9.4-4 |
| Kudu    | org.apache.kudu | kudu-client                 | 1.17.0.7.1.9.4-4 |
| Kudu    | org.apache.kudu | kudu-hive                   | 1.17.0.7.1.9.4-4 |
| Kudu    | org.apache.kudu | kudu-spark2-tools_2.11      | 1.17.0.7.1.9.4-4 |
| Kudu    | org.apache.kudu | kudu-spark2_2.11            | 1.17.0.7.1.9.4-4 |
| Kudu    | org.apache.kudu | kudu-spark3-tools_2.12      | 1.17.0.7.1.9.4-4 |
| Kudu    | org.apache.kudu | kudu-spark3_2.12            | 1.17.0.7.1.9.4-4 |
| Kudu    | org.apache.kudu | kudu-test-utils             | 1.17.0.7.1.9.4-4 |
| Livy    | org.apache.livy | livy-api                    | 0.7.2.7.1.9.4-4  |
| Livy    | org.apache.livy | livy-client-common          | 0.7.2.7.1.9.4-4  |
| Livy    | org.apache.livy | livy-client-http            | 0.7.2.7.1.9.4-4  |
| Livy    | org.apache.livy | livy-core_2.11              | 0.7.2.7.1.9.4-4  |

| Project | groupId           | artifactId                  | version          |
|---------|-------------------|-----------------------------|------------------|
| Livy    | org.apache.livy   | livy-examples               | 0.7.2.7.1.9.4-4  |
| Livy    | org.apache.livy   | livy-integration-test       | 0.7.2.7.1.9.4-4  |
| Livy    | org.apache.livy   | livy-repl_2.11              | 0.7.2.7.1.9.4-4  |
| Livy    | org.apache.livy   | livy-rsc                    | 0.7.2.7.1.9.4-4  |
| Livy    | org.apache.livy   | livy-scala-api_2.11         | 0.7.2.7.1.9.4-4  |
| Livy    | org.apache.livy   | livy-server                 | 0.7.2.7.1.9.4-4  |
| Livy    | org.apache.livy   | livy-test-lib               | 0.7.2.7.1.9.4-4  |
| Livy    | org.apache.livy   | livy-thriftserver           | 0.7.2.7.1.9.4-4  |
| Livy    | org.apache.livy   | livy-thriftserver-session   | 0.7.2.7.1.9.4-4  |
| Lucene  | org.apache.lucene | lucene-analyzers-common     | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-icu        | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-kuromoji   | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-morfologik | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-nori       | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-openslp    | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-phonetic   | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-smartcn    | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-analyzers-stempel    | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-backward-codecs      | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-benchmark            | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-classification       | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-codecs               | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-core                 | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-demo                 | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-expressions          | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-facet                | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-grouping             | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-highlighter          | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-join                 | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-memory               | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-misc                 | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-monitor              | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-queries              | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-queryparser          | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-replicator           | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-sandbox              | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-spatial-extras       | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-spatial3d            | 8.11.2.7.1.9.4-4 |
| Lucene  | org.apache.lucene | lucene-suggest              | 8.11.2.7.1.9.4-4 |

| Project | groupId           | artifactId                     | version          |
|---------|-------------------|--------------------------------|------------------|
| Lucene  | org.apache.lucene | lucene-test-framework          | 8.11.2.7.1.9.4-4 |
| Oozie   | org.apache.oozie  | oozie-client                   | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-core                     | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-distro                   | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-examples                 | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-api           | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-client        | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-server                   | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-sharelib-distcp          | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-sharelib-git             | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hcatalog        | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hive            | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hive2           | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-sharelib-oozie           | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-sharelib-spark           | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-sharelib-spark3          | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-sharelib-sqoop           | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-sharelib-streaming       | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-tools                    | 5.1.0.7.1.9.4-4  |
| Oozie   | org.apache.oozie  | oozie-zookeeper-security-tests | 5.1.0.7.1.9.4-4  |
| ORC     | org.apache.orc    | orc-core                       | 1.5.1.7.1.9.4-4  |
| ORC     | org.apache.orc    | orc-examples                   | 1.5.1.7.1.9.4-4  |
| ORC     | org.apache.orc    | orc-mapreduce                  | 1.5.1.7.1.9.4-4  |
| ORC     | org.apache.orc    | orc-shims                      | 1.5.1.7.1.9.4-4  |
| ORC     | org.apache.orc    | orc-tools                      | 1.5.1.7.1.9.4-4  |
| Ozone   | org.apache.ozone  | hdds-annotation-processing     | 1.4.0.7.1.9.4-4  |
| Ozone   | org.apache.ozone  | hdds-client                    | 1.4.0.7.1.9.4-4  |
| Ozone   | org.apache.ozone  | hdds-common                    | 1.4.0.7.1.9.4-4  |
| Ozone   | org.apache.ozone  | hdds-config                    | 1.4.0.7.1.9.4-4  |
| Ozone   | org.apache.ozone  | hdds-container-service         | 1.4.0.7.1.9.4-4  |
| Ozone   | org.apache.ozone  | hdds-docs                      | 1.4.0.7.1.9.4-4  |
| Ozone   | org.apache.ozone  | hdds-erasurencode              | 1.4.0.7.1.9.4-4  |
| Ozone   | org.apache.ozone  | hdds-hadoop-dependency-client  | 1.4.0.7.1.9.4-4  |
| Ozone   | org.apache.ozone  | hdds-hadoop-dependency-server  | 1.4.0.7.1.9.4-4  |
| Ozone   | org.apache.ozone  | hdds-hadoop-dependency-test    | 1.4.0.7.1.9.4-4  |
| Ozone   | org.apache.ozone  | hdds-interface-admin           | 1.4.0.7.1.9.4-4  |
| Ozone   | org.apache.ozone  | hdds-interface-client          | 1.4.0.7.1.9.4-4  |
| Ozone   | org.apache.ozone  | hdds-interface-server          | 1.4.0.7.1.9.4-4  |
| Ozone   | org.apache.ozone  | hdds-managed-rocksdbs          | 1.4.0.7.1.9.4-4  |

| Project | groupId            | artifactId                | version           |
|---------|--------------------|---------------------------|-------------------|
| Ozone   | org.apache.ozone   | hdds-rocks-native         | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | hdds-server-framework     | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | hdds-server-scm           | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | hdds-test-utils           | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | hdds-tools                | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | mini-chaos-tests          | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-client              | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-common              | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-csi                 | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-datanode            | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-dist                | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-filessystem         | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-filessystem-common  | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-filessystem-hadoop2 | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-filessystem-hadoop3 | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-filessystem-shaded  | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-httpfsgateway       | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-insight             | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-integration-test    | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-interface-client    | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-interface-storage   | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-manager             | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-network-tests       | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-recon               | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-reconcodegen        | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-s3-secret-store     | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-s3gateway           | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | ozone-tools               | 1.4.0.7.1.9.4-4   |
| Ozone   | org.apache.ozone   | rocksdb-checkpoint-differ | 1.4.0.7.1.9.4-4   |
| Parquet | org.apache.parquet | parquet-avro              | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-cascading         | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-cascading3        | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-column            | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-common            | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-encoding          | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-format-structures | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-generator         | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-hadoop            | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-hadoop-bundle     | 1.10.99.7.1.9.4-4 |

| Project | groupId            | artifactId                             | version           |
|---------|--------------------|----------------------------------------|-------------------|
| Parquet | org.apache.parquet | parquet-jackson                        | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-pig                            | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-pig-bundle                     | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-protobuf                       | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-scala_2.10                     | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-thrift                         | 1.10.99.7.1.9.4-4 |
| Parquet | org.apache.parquet | parquet-tools                          | 1.10.99.7.1.9.4-4 |
| Phoenix | org.apache.phoenix | phoenix-client-embedded-hbase-2.4      | 5.1.1.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-client-hbase-2.4               | 5.1.1.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-connectors-phoenix5-compatible | 6.0.0.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-core                           | 5.1.1.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.1.6         | 5.1.1.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.2.5         | 5.1.1.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.3.0         | 5.1.1.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.4.0         | 5.1.1.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.4.1         | 5.1.1.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-pherf                          | 5.1.1.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-queryserver                    | 6.0.0.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-client             | 6.0.0.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-it                 | 6.0.0.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-load-balancer      | 6.0.0.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-orchestrator       | 6.0.0.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-server-hbase-2.4               | 5.1.1.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix-tracing-webapp                 | 5.1.1.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix5-hive                          | 6.0.0.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix5-hive-shaded                   | 6.0.0.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix5-spark                         | 6.0.0.7.1.9.4-4   |
| Phoenix | org.apache.phoenix | phoenix5-spark-shaded                  | 6.0.0.7.1.9.4-4   |
| Ranger  | org.apache.ranger  | conditions-enrichers                   | 2.4.0.7.1.9.4-4   |
| Ranger  | org.apache.ranger  | credentialbuilder                      | 2.4.0.7.1.9.4-4   |
| Ranger  | org.apache.ranger  | embeddedwebserver                      | 2.4.0.7.1.9.4-4   |
| Ranger  | org.apache.ranger  | jisql                                  | 2.4.0.7.1.9.4-4   |
| Ranger  | org.apache.ranger  | ldapconfigcheck                        | 2.4.0.7.1.9.4-4   |
| Ranger  | org.apache.ranger  | ranger-adls-plugin                     | 2.4.0.7.1.9.4-4   |
| Ranger  | org.apache.ranger  | ranger-atlas-plugin                    | 2.4.0.7.1.9.4-4   |
| Ranger  | org.apache.ranger  | ranger-atlas-plugin-shim               | 2.4.0.7.1.9.4-4   |
| Ranger  | org.apache.ranger  | ranger-authn                           | 2.4.0.7.1.9.4-4   |
| Ranger  | org.apache.ranger  | ranger-common-ha                       | 2.4.0.7.1.9.4-4   |
| Ranger  | org.apache.ranger  | ranger-distro                          | 2.4.0.7.1.9.4-4   |

| Project | groupId           | artifactId                  | version         |
|---------|-------------------|-----------------------------|-----------------|
| Ranger  | org.apache.ranger | ranger-examples-distro      | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-hbase-plugin         | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-hbase-plugin-shim    | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-hdfs-plugin          | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-hdfs-plugin-shim     | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-hive-plugin          | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-hive-plugin-shim     | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-intg                 | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-kafka-connect-plugin | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-kafka-plugin         | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-kafka-plugin-shim    | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-kms                  | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-kms-plugin           | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-kms-plugin-shim      | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-knox-plugin          | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-knox-plugin-shim     | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-kudu-plugin          | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-kylin-plugin         | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-kylin-plugin-shim    | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-metrics              | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-nifi-plugin          | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-nifi-registry-plugin | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-ozone-plugin         | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-ozone-plugin-shim    | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-plugin-classloader   | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-plugins-audit        | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-plugins-common       | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-plugins-cred         | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-plugins-installer    | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-policymigration      | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-raz-adls             | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-raz-chained-plugins  | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-raz-hook-abfs        | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-raz-intg             | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-raz-processor        | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-rms-common           | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-rms-hive             | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-rms-plugins-common   | 2.4.0.7.1.9.4-4 |
| Ranger  | org.apache.ranger | ranger-rms-webapp           | 2.4.0.7.1.9.4-4 |

| Project | groupId           | artifactId                     | version          |
|---------|-------------------|--------------------------------|------------------|
| Ranger  | org.apache.ranger | ranger-sampleapp-plugin        | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | ranger-schema-registry-plugin  | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | ranger-solr-plugin             | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | ranger-solr-plugin-shim        | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin            | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin-shim       | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | ranger-storm-plugin            | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | ranger-storm-plugin-shim       | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | ranger-tagsync                 | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | ranger-tools                   | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | ranger-util                    | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | ranger-yarn-plugin             | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | ranger-yarn-plugin-shim        | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | sample-client                  | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | sampleapp                      | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | shaded-raz-hook-abfs           | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | ugsync-util                    | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | unixauthclient                 | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | unixauthservice                | 2.4.0.7.1.9.4-4  |
| Ranger  | org.apache.ranger | unixusersync                   | 2.4.0.7.1.9.4-4  |
| Solr    | org.apache.solr   | solr-analysis-extras           | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-analytics                 | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-cell                      | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-core                      | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-dataimporthandler         | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-dataimporthandler-extras  | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-gcs-repository            | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-jaegertracer-configurator | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-langid                    | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-ltr                       | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-prometheus-exporter       | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-s3-repository             | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-security-util             | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-solrj                     | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-test-framework            | 8.11.2.7.1.9.4-4 |
| Solr    | org.apache.solr   | solr-velocity                  | 8.11.2.7.1.9.4-4 |
| Spark   | org.apache.spark  | spark-avro_2.11                | 2.4.8.7.1.9.4-4  |
| Spark   | org.apache.spark  | spark-catalyst_2.11            | 2.4.8.7.1.9.4-4  |
| Spark   | org.apache.spark  | spark-core_2.11                | 2.4.8.7.1.9.4-4  |

| Project | groupId          | artifactId                               | version         |
|---------|------------------|------------------------------------------|-----------------|
| Spark   | org.apache.spark | spark-graphx_2.11                        | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-hadoop-cloud_2.11                  | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-hive_2.11                          | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-kubernetes_2.11                    | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-kvstore_2.11                       | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-launcher_2.11                      | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-mllib-local_2.11                   | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-mllib_2.11                         | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-network-common_2.11                | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-network-shuffle_2.11               | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-network-yarn_2.11                  | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-repl_2.11                          | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-sketch_2.11                        | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-sql-kafka-0-10_2.11                | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-sql_2.11                           | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-streaming-kafka-0-10-assembly_2.11 | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-streaming-kafka-0-10_2.11          | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-streaming_2.11                     | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-tags_2.11                          | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-token-provider-kafka-0-10_2.11     | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-unsafe_2.11                        | 2.4.8.7.1.9.4-4 |
| Spark   | org.apache.spark | spark-yarn_2.11                          | 2.4.8.7.1.9.4-4 |
| Sqoop   | org.apache.sqoop | sqoop                                    | 1.4.7.7.1.9.4-4 |
| Sqoop   | org.apache.sqoop | sqoop-test                               | 1.4.7.7.1.9.4-4 |
| Tez     | org.apache.tez   | hadoop-shim                              | 0.9.1.7.1.9.4-4 |
| Tez     | org.apache.tez   | hadoop-shim-2.8                          | 0.9.1.7.1.9.4-4 |
| Tez     | org.apache.tez   | tez-api                                  | 0.9.1.7.1.9.4-4 |
| Tez     | org.apache.tez   | tez-aux-services                         | 0.9.1.7.1.9.4-4 |
| Tez     | org.apache.tez   | tez-common                               | 0.9.1.7.1.9.4-4 |
| Tez     | org.apache.tez   | tez-dag                                  | 0.9.1.7.1.9.4-4 |
| Tez     | org.apache.tez   | tez-examples                             | 0.9.1.7.1.9.4-4 |
| Tez     | org.apache.tez   | tez-ext-service-tests                    | 0.9.1.7.1.9.4-4 |
| Tez     | org.apache.tez   | tez-history-parser                       | 0.9.1.7.1.9.4-4 |
| Tez     | org.apache.tez   | tez-javadoc-tools                        | 0.9.1.7.1.9.4-4 |
| Tez     | org.apache.tez   | tez-job-analyzer                         | 0.9.1.7.1.9.4-4 |
| Tez     | org.apache.tez   | tez-mapreduce                            | 0.9.1.7.1.9.4-4 |
| Tez     | org.apache.tez   | tez-protobuf-history-plugin              | 0.9.1.7.1.9.4-4 |
| Tez     | org.apache.tez   | tez-runtime-internals                    | 0.9.1.7.1.9.4-4 |
| Tez     | org.apache.tez   | tez-runtime-library                      | 0.9.1.7.1.9.4-4 |



| Project   | groupId              | artifactId                          | version         |
|-----------|----------------------|-------------------------------------|-----------------|
| Tez       | org.apache.tez       | tez-tests                           | 0.9.1.7.1.9.4-4 |
| Tez       | org.apache.tez       | tez-yarn-timeline-cache-plugin      | 0.9.1.7.1.9.4-4 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history           | 0.9.1.7.1.9.4-4 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history-with-acls | 0.9.1.7.1.9.4-4 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history-with-fs   | 0.9.1.7.1.9.4-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-angular                    | 0.8.2.7.1.9.4-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-display                    | 0.8.2.7.1.9.4-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-interpreter                | 0.8.2.7.1.9.4-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-jdbc                       | 0.8.2.7.1.9.4-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-jupyter                    | 0.8.2.7.1.9.4-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-livy                       | 0.8.2.7.1.9.4-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-markdown                   | 0.8.2.7.1.9.4-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-server                     | 0.8.2.7.1.9.4-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-shell                      | 0.8.2.7.1.9.4-4 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-zengine                    | 0.8.2.7.1.9.4-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper                           | 3.8.1.7.1.9.4-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-fatjar            | 3.8.1.7.1.9.4-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-loggraph          | 3.8.1.7.1.9.4-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-rest              | 3.8.1.7.1.9.4-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-zooinspector      | 3.8.1.7.1.9.4-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-it                        | 3.8.1.7.1.9.4-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-jute                      | 3.8.1.7.1.9.4-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-prometheus-metrics        | 3.8.1.7.1.9.4-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-election          | 3.8.1.7.1.9.4-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-lock              | 3.8.1.7.1.9.4-4 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-queue             | 3.8.1.7.1.9.4-4 |

### Runtime 7.1.9.6-3

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          | 3.0.0.7.1.9.6-3 |
| Atlas   | org.apache.atlas | atlas-aws-s3-bridge          | 3.0.0.7.1.9.6-3 |
| Atlas   | org.apache.atlas | atlas-azure-adls-bridge      | 3.0.0.7.1.9.6-3 |
| Atlas   | org.apache.atlas | atlas-classification-updater | 3.0.0.7.1.9.6-3 |
| Atlas   | org.apache.atlas | atlas-client-common          | 3.0.0.7.1.9.6-3 |
| Atlas   | org.apache.atlas | atlas-client-v1              | 3.0.0.7.1.9.6-3 |
| Atlas   | org.apache.atlas | atlas-client-v2              | 3.0.0.7.1.9.6-3 |
| Atlas   | org.apache.atlas | atlas-common                 | 3.0.0.7.1.9.6-3 |
| Atlas   | org.apache.atlas | atlas-distro                 | 3.0.0.7.1.9.6-3 |
| Atlas   | org.apache.atlas | atlas-docs                   | 3.0.0.7.1.9.6-3 |

| Project | groupId          | artifactId               | version          |
|---------|------------------|--------------------------|------------------|
| Atlas   | org.apache.atlas | atlas-graphdb-api        | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | atlas-graphdb-common     | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | atlas-graphdb-janus      | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | atlas-hdfs-bridge        | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | atlas-index-repair-tool  | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | atlas-intg               | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | atlas-janusgraph-hbase2  | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | atlas-notification       | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | atlas-plugin-classloader | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | atlas-repository         | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | atlas-server-api         | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | atlas-testtools          | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | hbase-bridge             | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | hbase-bridge-shim        | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | hbase-testing-util       | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | hdfs-model               | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | hive-bridge              | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | hive-bridge-shim         | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | impala-bridge            | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | impala-bridge-shim       | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | impala-hook-api          | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | kafka-bridge             | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | kafka-bridge-shim        | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | navigator-to-atlas       | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | sample-app               | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | sqoop-bridge             | 3.0.0.7.1.9.6-3  |
| Atlas   | org.apache.atlas | sqoop-bridge-shim        | 3.0.0.7.1.9.6-3  |
| Avro    | org.apache.avro  | avro                     | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro  | avro-android             | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro  | avro-codegen-test        | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro  | avro-compiler            | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro  | avro-grpc                | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro  | avro-ipc                 | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro  | avro-ipc-jetty           | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro  | avro-ipc-netty           | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro  | avro-mapred              | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro  | avro-maven-plugin        | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro  | avro-perf                | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro  | avro-protobuf            | 1.11.1.7.1.9.6-3 |

| Project | groupId            | artifactId                      | version          |
|---------|--------------------|---------------------------------|------------------|
| Avro    | org.apache.avro    | avro-service-archetype          | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro    | avro-test-custom-conversions    | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro    | avro-thrift                     | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro    | avro-tools                      | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro    | trevni-avro                     | 1.11.1.7.1.9.6-3 |
| Avro    | org.apache.avro    | trevni-core                     | 1.11.1.7.1.9.6-3 |
| Calcite | org.apache.calcite | calcite-babel                   | 1.19.0.7.1.9.6-3 |
| Calcite | org.apache.calcite | calcite-core                    | 1.19.0.7.1.9.6-3 |
| Calcite | org.apache.calcite | calcite-druid                   | 1.19.0.7.1.9.6-3 |
| Calcite | org.apache.calcite | calcite-linq4j                  | 1.19.0.7.1.9.6-3 |
| Calcite | org.apache.calcite | calcite-server                  | 1.19.0.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop  | hadoop-aliyun                   | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-annotations              | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-archive-logs             | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-archives                 | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-assemblies               | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-auth                     | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-aws                      | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-azure                    | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-azure-datalake           | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-build-tools              | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-client                   | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-client-api               | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-client-integration-tests | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-client-minicluster       | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-client-runtime           | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-cloud-storage            | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-common                   | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-datajoin                 | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-distcp                   | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-extras                   | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-fs2img                   | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-gridmix                  | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-hdfs                     | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-hdfs-client              | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-hdfs-httpfs              | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-hdfs-native-client       | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-hdfs-nfs                 | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop  | hadoop-hdfs-rbf                 | 3.1.1.7.1.9.6-3  |

| Project | groupId           | artifactId                                      | version         |
|---------|-------------------|-------------------------------------------------|-----------------|
| Hadoop  | org.apache.hadoop | hadoop-kafka                                    | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-kms                                      | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-app                     | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-common                  | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-core                    | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-hs                      | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-hs-plugins              | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-jobclient               | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-nativetask              | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-shuffle                 | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-uploader                | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-examples                       | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-maven-plugins                            | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-minicluster                              | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-minikdc                                  | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-nfs                                      | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-openstack                                | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-resourceestimator                        | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-rumen                                    | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-sls                                      | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-streaming                                | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-tools-dist                               | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-api                                 | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-distributedshell       | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-unmanaged-am-launcher  | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-client                              | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-common                              | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-registry                            | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-applicationhistoryservice    | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-common                       | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-nodemanager                  | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-resourcemanager              | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-router                       | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-sharedcachemanager           | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-tests                        | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timeline-pluginstorage       | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice              | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-client | 3.1.1.7.1.9.6-3 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-common | 3.1.1.7.1.9.6-3 |

| Project | groupId           | artifactId                                        | version          |
|---------|-------------------|---------------------------------------------------|------------------|
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-server-2 | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-tests    | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-web-proxy                      | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-api                          | 3.1.1.7.1.9.6-3  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-core                         | 3.1.1.7.1.9.6-3  |
| HBase   | org.apache.hbase  | hbase-annotations                                 | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-asyncfs                                     | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-checkstyle                                  | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-client                                      | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-client-project                              | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-common                                      | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-endpoint                                    | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-examples                                    | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-external-blockcache                         | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-hadoop-compat                               | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-hadoop2-compat                              | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-hbtop                                       | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-http                                        | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-it                                          | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-logging                                     | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-mapreduce                                   | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-metrics                                     | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-metrics-api                                 | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-procedure                                   | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-protocol                                    | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-protocol-shaded                             | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-replication                                 | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-resource-bundle                             | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-rest                                        | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-rsgroup                                     | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-server                                      | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-shaded-client                               | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-shaded-client-byo-hadoop                    | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-shaded-client-project                       | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-shaded-mapreduce                            | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-shaded-testing-util                         | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-shaded-testing-util-tester                  | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-shell                                       | 2.4.17.7.1.9.6-3 |
| HBase   | org.apache.hbase  | hbase-testing-util                                | 2.4.17.7.1.9.6-3 |

| Project | groupId          | artifactId                | version            |
|---------|------------------|---------------------------|--------------------|
| HBase   | org.apache.hbase | hbase-thrift              | 2.4.17.7.1.9.6-3   |
| HBase   | org.apache.hbase | hbase-zookeeper           | 2.4.17.7.1.9.6-3   |
| Hive    | org.apache.hive  | hive-beeline              | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-blobstore            | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-classification       | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-cli                  | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-common               | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-contrib              | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-exec                 | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-hbase-handler        | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-hcatalog-it-unit     | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-hplsql               | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-it-custom-serde      | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-it-minikdc           | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-it-qfile             | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-it-qfile-kudu        | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-it-test-serde        | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-it-unit              | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-it-unit-hadoop2      | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-it-util              | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-jdbc                 | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-jdbc-handler         | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-jmh                  | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-kryo-registrator     | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-kudu-handler         | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-llap-client          | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-llap-common          | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-llap-ext-client      | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-llap-server          | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-llap-tez             | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-metastore            | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-pre-upgrade          | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-serde                | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-service              | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-service-rpc          | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-shims                | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-spark-client         | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-standalone-metastore | 3.1.3000.7.1.9.6-3 |
| Hive    | org.apache.hive  | hive-storage-api          | 3.1.3000.7.1.9.6-3 |

| Project                  | groupId              | artifactId                               | version            |
|--------------------------|----------------------|------------------------------------------|--------------------|
| Hive                     | org.apache.hive      | hive-streaming                           | 3.1.3000.7.1.9.6-3 |
| Hive                     | org.apache.hive      | hive-testutils                           | 3.1.3000.7.1.9.6-3 |
| Hive                     | org.apache.hive      | hive-vector-code-gen                     | 3.1.3000.7.1.9.6-3 |
| Hive                     | org.apache.hive      | kafka-handler                            | 3.1.3000.7.1.9.6-3 |
| Hive Warehouse Connector | com.hortonworks.hive | hive-warehouse-connector_2.11            | 1.0.0.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | ci                                       | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | connect                                  | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | connect-api                              | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | connect-basic-auth-extension             | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | connect-cloudera-authorization-extension | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | connect-cloudera-common                  | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | connect-cloudera-secret-storage          | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | connect-cloudera-security-policies       | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | connect-file                             | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | connect-json                             | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | connect-mirror                           | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | connect-mirror-client                    | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | connect-runtime                          | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | connect-transforms                       | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | generator                                | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-clients                            | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-cloudera-metrics-reporter_2.12     | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-cloudera-metrics-reporter_2.13     | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-cloudera-plugins                   | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-examples                           | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-group-coordinator                  | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-log4j-appender                     | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-metadata                           | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-raft                               | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-server-common                      | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-shell                              | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-storage                            | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-storage-api                        | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-streams                            | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-streams-examples                   | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-streams-scala_2.12                 | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-streams-scala_2.13                 | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-streams-test-utils                 | 3.4.1.7.1.9.6-3    |
| Kafka                    | org.apache.kafka     | kafka-streams-upgrade-system-tests-0100  | 3.4.1.7.1.9.6-3    |

| Project | groupId          | artifactId                                        | version         |
|---------|------------------|---------------------------------------------------|-----------------|
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0101           | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0102           | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0110           | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-10             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-11             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-20             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-21             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-22             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-23             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-24             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-25             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-26             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-27             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-28             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-30             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-31             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-32             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-33             | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka-tools                                       | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka_2.12                                        | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | kafka_2.13                                        | 3.4.1.7.1.9.6-3 |
| Kafka   | org.apache.kafka | trogdor                                           | 3.4.1.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-adapter                                   | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-admin-ui                                  | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-applications                              | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-cloud-bindings                            | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-demo-ldap                                 | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-demo-ldap-launcher                        | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-discovery-ambari                          | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-discovery-cm                              | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-docker                                    | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-i18n                                      | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-i18n-logging-log4j                        | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-i18n-logging-sl4j                         | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-performance-test                          | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-provider-ha                               | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-common        | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-concat        | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-hadoop-groups | 1.3.0.7.1.9.6-3 |



| Project | groupId         | artifactId                                        | version         |
|---------|-----------------|---------------------------------------------------|-----------------|
| Knox    | org.apache.knox | gateway-provider-identity-assertion-pseudo        | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-regex         | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-switchcase    | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-jersey                           | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-rewrite                          | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-rewrite-common                   | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-hostmap-static      | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-inbound-query-param | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-service-registry    | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-rewrite-step-encrypt-uri         | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-rewrite-step-secure-query        | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-security-authc-anon              | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-security-authz-acls              | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-security-authz-composite         | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-security-clientcert              | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-security-hadoopauth              | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-security-jwt                     | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-security-pac4j                   | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-security-preauth                 | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-security-shiro                   | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-provider-security-webappsec               | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-release                                   | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-server                                    | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-server-launcher                           | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-server-xforwarded-filter                  | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-service-admin                             | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-service-as                                | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-service-definitions                       | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-service-hashicorp-vault                   | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-service-hbase                             | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-service-health                            | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-service-hive                              | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-service-idbroker                          | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-service-impala                            | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-service-jkg                               | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-service-knoxsso                           | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-service-knoxssout                         | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-service-knoxtoken                         | 1.3.0.7.1.9.6-3 |
| Knox    | org.apache.knox | gateway-service-livy                              | 1.3.0.7.1.9.6-3 |

| Project | groupId         | artifactId                    | version          |
|---------|-----------------|-------------------------------|------------------|
| Knox    | org.apache.knox | gateway-service-metadata      | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-service-nifi          | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-service-nifi-registry | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-service-remoteconfig  | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-service-rm            | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-service-session       | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-service-storm         | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-service-test          | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-service-tgs           | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-service-vault         | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-service-webhdfs       | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-shell                 | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-shell-launcher        | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-shell-release         | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-shell-samples         | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-spi                   | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-test                  | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-test-idbroker         | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-test-release-utils    | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-test-utils            | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-topology-hadoop-xml   | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-topology-simple       | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-util-common           | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-util-configinjector   | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-util-launcher         | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | gateway-util-urltemplate      | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | hadoop-examples               | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | knox-cli-launcher             | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | knox-homepage-ui              | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | knox-token-generation-ui      | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | knox-token-management-ui      | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | webhdfs-kerb-test             | 1.3.0.7.1.9.6-3  |
| Knox    | org.apache.knox | webhdfs-test                  | 1.3.0.7.1.9.6-3  |
| Kudu    | org.apache.kudu | kudu-backup-tools             | 1.17.0.7.1.9.6-3 |
| Kudu    | org.apache.kudu | kudu-backup2_2.11             | 1.17.0.7.1.9.6-3 |
| Kudu    | org.apache.kudu | kudu-backup3_2.12             | 1.17.0.7.1.9.6-3 |
| Kudu    | org.apache.kudu | kudu-client                   | 1.17.0.7.1.9.6-3 |
| Kudu    | org.apache.kudu | kudu-hive                     | 1.17.0.7.1.9.6-3 |
| Kudu    | org.apache.kudu | kudu-spark2-tools_2.11        | 1.17.0.7.1.9.6-3 |

| Project | groupId           | artifactId                  | version          |
|---------|-------------------|-----------------------------|------------------|
| Kudu    | org.apache.kudu   | kudu-spark2_2.11            | 1.17.0.7.1.9.6-3 |
| Kudu    | org.apache.kudu   | kudu-spark3-tools_2.12      | 1.17.0.7.1.9.6-3 |
| Kudu    | org.apache.kudu   | kudu-spark3_2.12            | 1.17.0.7.1.9.6-3 |
| Kudu    | org.apache.kudu   | kudu-test-utils             | 1.17.0.7.1.9.6-3 |
| Livy    | org.apache.livy   | livy-api                    | 0.7.2.7.1.9.6-3  |
| Livy    | org.apache.livy   | livy-client-common          | 0.7.2.7.1.9.6-3  |
| Livy    | org.apache.livy   | livy-client-http            | 0.7.2.7.1.9.6-3  |
| Livy    | org.apache.livy   | livy-core_2.11              | 0.7.2.7.1.9.6-3  |
| Livy    | org.apache.livy   | livy-examples               | 0.7.2.7.1.9.6-3  |
| Livy    | org.apache.livy   | livy-integration-test       | 0.7.2.7.1.9.6-3  |
| Livy    | org.apache.livy   | livy-repl_2.11              | 0.7.2.7.1.9.6-3  |
| Livy    | org.apache.livy   | livy-rsc                    | 0.7.2.7.1.9.6-3  |
| Livy    | org.apache.livy   | livy-scala-api_2.11         | 0.7.2.7.1.9.6-3  |
| Livy    | org.apache.livy   | livy-server                 | 0.7.2.7.1.9.6-3  |
| Livy    | org.apache.livy   | livy-test-lib               | 0.7.2.7.1.9.6-3  |
| Livy    | org.apache.livy   | livy-thriftserver           | 0.7.2.7.1.9.6-3  |
| Livy    | org.apache.livy   | livy-thriftserver-session   | 0.7.2.7.1.9.6-3  |
| Lucene  | org.apache.lucene | lucene-analyzers-common     | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-analyzers-icu        | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-analyzers-kuromoji   | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-analyzers-morfologik | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-analyzers-nori       | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-analyzers-openslp    | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-analyzers-phonetic   | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-analyzers-smartcn    | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-analyzers-stempel    | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-backward-codecs      | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-benchmark            | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-classification       | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-codecs               | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-core                 | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-demo                 | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-expressions          | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-facet                | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-grouping             | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-highlighter          | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-join                 | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-memory               | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-misc                 | 8.11.2.7.1.9.6-3 |

| Project | groupId           | artifactId                     | version          |
|---------|-------------------|--------------------------------|------------------|
| Lucene  | org.apache.lucene | lucene-monitor                 | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-queries                 | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-queryparser             | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-replicator              | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-sandbox                 | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-spatial-extras          | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-spatial3d               | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-suggest                 | 8.11.2.7.1.9.6-3 |
| Lucene  | org.apache.lucene | lucene-test-framework          | 8.11.2.7.1.9.6-3 |
| Oozie   | org.apache.oozie  | oozie-client                   | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-core                     | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-distro                   | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-examples                 | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-api           | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-client        | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-server                   | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-sharelib-distcp          | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-sharelib-git             | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hcatalog        | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hive            | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hive2           | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-sharelib-oozie           | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-sharelib-spark           | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-sharelib-spark3          | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-sharelib-sqoop           | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-sharelib-streaming       | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-tools                    | 5.1.0.7.1.9.6-3  |
| Oozie   | org.apache.oozie  | oozie-zookeeper-security-tests | 5.1.0.7.1.9.6-3  |
| ORC     | org.apache.orc    | orc-core                       | 1.5.1.7.1.9.6-3  |
| ORC     | org.apache.orc    | orc-examples                   | 1.5.1.7.1.9.6-3  |
| ORC     | org.apache.orc    | orc-mapreduce                  | 1.5.1.7.1.9.6-3  |
| ORC     | org.apache.orc    | orc-shims                      | 1.5.1.7.1.9.6-3  |
| ORC     | org.apache.orc    | orc-tools                      | 1.5.1.7.1.9.6-3  |
| Ozone   | org.apache.ozone  | hdds-annotation-processing     | 1.4.0.7.1.9.6-3  |
| Ozone   | org.apache.ozone  | hdds-client                    | 1.4.0.7.1.9.6-3  |
| Ozone   | org.apache.ozone  | hdds-common                    | 1.4.0.7.1.9.6-3  |
| Ozone   | org.apache.ozone  | hdds-config                    | 1.4.0.7.1.9.6-3  |
| Ozone   | org.apache.ozone  | hdds-container-service         | 1.4.0.7.1.9.6-3  |
| Ozone   | org.apache.ozone  | hdds-docs                      | 1.4.0.7.1.9.6-3  |

| Project | groupId            | artifactId                    | version           |
|---------|--------------------|-------------------------------|-------------------|
| Ozone   | org.apache.ozone   | hdds-erasurecode              | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | hdds-hadoop-dependency-client | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | hdds-hadoop-dependency-server | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | hdds-hadoop-dependency-test   | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | hdds-interface-admin          | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | hdds-interface-client         | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | hdds-interface-server         | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | hdds-managed-rocksdb          | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | hdds-rocks-native             | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | hdds-server-framework         | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | hdds-server-scm               | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | hdds-test-utils               | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | hdds-tools                    | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | mini-chaos-tests              | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-client                  | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-common                  | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-csi                     | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-datanode                | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-dist                    | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-filessystem             | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-filessystem-common      | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-filessystem-hadoop2     | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-filessystem-hadoop3     | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-filessystem-shaded      | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-httpsgateway            | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-insight                 | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-integration-test        | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-interface-client        | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-interface-storage       | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-manager                 | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-network-tests           | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-recon                   | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-reconcodegen            | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-s3-secret-store         | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-s3gateway               | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | ozone-tools                   | 1.4.0.7.1.9.6-3   |
| Ozone   | org.apache.ozone   | rocksdb-checkpoint-differ     | 1.4.0.7.1.9.6-3   |
| Parquet | org.apache.parquet | parquet-avro                  | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-cascading             | 1.10.99.7.1.9.6-3 |

| Project | groupId            | artifactId                             | version           |
|---------|--------------------|----------------------------------------|-------------------|
| Parquet | org.apache.parquet | parquet-cascading3                     | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-column                         | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-common                         | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-encoding                       | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-format-structures              | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-generator                      | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-hadoop                         | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-hadoop-bundle                  | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-jackson                        | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-pig                            | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-pig-bundle                     | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-protobuf                       | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-scala_2.10                     | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-thrift                         | 1.10.99.7.1.9.6-3 |
| Parquet | org.apache.parquet | parquet-tools                          | 1.10.99.7.1.9.6-3 |
| Phoenix | org.apache.phoenix | phoenix-client-embedded-hbase-2.4      | 5.1.1.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-client-hbase-2.4               | 5.1.1.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-connectors-phoenix5-compatible | 6.0.0.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-core                           | 5.1.1.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.1.6         | 5.1.1.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.2.5         | 5.1.1.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.3.0         | 5.1.1.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.4.0         | 5.1.1.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.4.1         | 5.1.1.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-pherf                          | 5.1.1.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-queryserver                    | 6.0.0.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-client             | 6.0.0.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-it                 | 6.0.0.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-load-balancer      | 6.0.0.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-orchestrator       | 6.0.0.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-server-hbase-2.4               | 5.1.1.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix-tracing-webapp                 | 5.1.1.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix5-hive                          | 6.0.0.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix5-hive-shaded                   | 6.0.0.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix5-spark                         | 6.0.0.7.1.9.6-3   |
| Phoenix | org.apache.phoenix | phoenix5-spark-shaded                  | 6.0.0.7.1.9.6-3   |
| Ranger  | org.apache.ranger  | conditions-enrichers                   | 2.4.0.7.1.9.6-3   |
| Ranger  | org.apache.ranger  | credentialbuilder                      | 2.4.0.7.1.9.6-3   |
| Ranger  | org.apache.ranger  | embeddedwebserver                      | 2.4.0.7.1.9.6-3   |

| Project | groupId           | artifactId                  | version         |
|---------|-------------------|-----------------------------|-----------------|
| Ranger  | org.apache.ranger | jisql                       | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ldapconfigcheck             | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-adls-plugin          | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-atlas-plugin         | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-atlas-plugin-shim    | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-authn                | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-common-ha            | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-distro               | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-examples-distro      | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-hbase-plugin         | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-hbase-plugin-shim    | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-hdfs-plugin          | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-hdfs-plugin-shim     | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-hive-plugin          | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-hive-plugin-shim     | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-intg                 | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-kafka-connect-plugin | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-kafka-plugin         | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-kafka-plugin-shim    | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-kms                  | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-kms-plugin           | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-kms-plugin-shim      | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-knox-plugin          | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-knox-plugin-shim     | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-kudu-plugin          | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-kylin-plugin         | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-kylin-plugin-shim    | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-metrics              | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-nifi-plugin          | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-nifi-registry-plugin | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-ozone-plugin         | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-ozone-plugin-shim    | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-plugin-classloader   | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-plugins-audit        | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-plugins-common       | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-plugins-cred         | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-plugins-installer    | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-policymigration      | 2.4.0.7.1.9.6-3 |
| Ranger  | org.apache.ranger | ranger-raz-adls             | 2.4.0.7.1.9.6-3 |

| Project | groupId           | artifactId                     | version          |
|---------|-------------------|--------------------------------|------------------|
| Ranger  | org.apache.ranger | ranger-raz-chained-plugins     | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-raz-hook-abfs           | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-raz-intg                | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-raz-processor           | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-rms-common              | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-rms-hive                | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-rms-plugins-common      | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-rms-webapp              | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-sampleapp-plugin        | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-schema-registry-plugin  | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-solr-plugin             | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-solr-plugin-shim        | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin            | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin-shim       | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-storm-plugin            | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-storm-plugin-shim       | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-tagsync                 | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-tools                   | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-util                    | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-yarn-plugin             | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ranger-yarn-plugin-shim        | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | sample-client                  | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | sampleapp                      | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | shaded-raz-hook-abfs           | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | ugsync-util                    | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | unixauthclient                 | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | unixauthservice                | 2.4.0.7.1.9.6-3  |
| Ranger  | org.apache.ranger | unixusersync                   | 2.4.0.7.1.9.6-3  |
| Solr    | org.apache.solr   | solr-analysis-extras           | 8.11.2.7.1.9.6-3 |
| Solr    | org.apache.solr   | solr-analytics                 | 8.11.2.7.1.9.6-3 |
| Solr    | org.apache.solr   | solr-cell                      | 8.11.2.7.1.9.6-3 |
| Solr    | org.apache.solr   | solr-core                      | 8.11.2.7.1.9.6-3 |
| Solr    | org.apache.solr   | solr-dataimporthandler         | 8.11.2.7.1.9.6-3 |
| Solr    | org.apache.solr   | solr-dataimporthandler-extras  | 8.11.2.7.1.9.6-3 |
| Solr    | org.apache.solr   | solr-gcs-repository            | 8.11.2.7.1.9.6-3 |
| Solr    | org.apache.solr   | solr-jaegertracer-configurator | 8.11.2.7.1.9.6-3 |
| Solr    | org.apache.solr   | solr-langid                    | 8.11.2.7.1.9.6-3 |
| Solr    | org.apache.solr   | solr-ltr                       | 8.11.2.7.1.9.6-3 |
| Solr    | org.apache.solr   | solr-prometheus-exporter       | 8.11.2.7.1.9.6-3 |



| Project | groupId          | artifactId                               | version          |
|---------|------------------|------------------------------------------|------------------|
| Solr    | org.apache.solr  | solr-s3-repository                       | 8.11.2.7.1.9.6-3 |
| Solr    | org.apache.solr  | solr-security-util                       | 8.11.2.7.1.9.6-3 |
| Solr    | org.apache.solr  | solr-solrj                               | 8.11.2.7.1.9.6-3 |
| Solr    | org.apache.solr  | solr-test-framework                      | 8.11.2.7.1.9.6-3 |
| Solr    | org.apache.solr  | solr-velocity                            | 8.11.2.7.1.9.6-3 |
| Spark   | org.apache.spark | spark-avro_2.11                          | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-catalyst_2.11                      | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-core_2.11                          | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-graphx_2.11                        | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-hadoop-cloud_2.11                  | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-hive_2.11                          | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-kubernetes_2.11                    | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-kvstore_2.11                       | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-launcher_2.11                      | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-mllib-local_2.11                   | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-mllib_2.11                         | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-network-common_2.11                | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-network-shuffle_2.11               | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-network-yarn_2.11                  | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-repl_2.11                          | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-sketch_2.11                        | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-sql-kafka-0-10_2.11                | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-sql_2.11                           | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-streaming-kafka-0-10-assembly_2.11 | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-streaming-kafka-0-10_2.11          | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-streaming_2.11                     | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-tags_2.11                          | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-token-provider-kafka-0-10_2.11     | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-unsafe_2.11                        | 2.4.8.7.1.9.6-3  |
| Spark   | org.apache.spark | spark-yarn_2.11                          | 2.4.8.7.1.9.6-3  |
| Sqoop   | org.apache.sqoop | sqoop                                    | 1.4.7.7.1.9.6-3  |
| Sqoop   | org.apache.sqoop | sqoop-test                               | 1.4.7.7.1.9.6-3  |
| Tez     | org.apache.tez   | hadoop-shim                              | 0.9.1.7.1.9.6-3  |
| Tez     | org.apache.tez   | hadoop-shim-2.8                          | 0.9.1.7.1.9.6-3  |
| Tez     | org.apache.tez   | tez-api                                  | 0.9.1.7.1.9.6-3  |
| Tez     | org.apache.tez   | tez-aux-services                         | 0.9.1.7.1.9.6-3  |
| Tez     | org.apache.tez   | tez-common                               | 0.9.1.7.1.9.6-3  |
| Tez     | org.apache.tez   | tez-dag                                  | 0.9.1.7.1.9.6-3  |
| Tez     | org.apache.tez   | tez-examples                             | 0.9.1.7.1.9.6-3  |

| Project   | groupId              | artifactId                          | version         |
|-----------|----------------------|-------------------------------------|-----------------|
| Tez       | org.apache.tez       | tez-ext-service-tests               | 0.9.1.7.1.9.6-3 |
| Tez       | org.apache.tez       | tez-history-parser                  | 0.9.1.7.1.9.6-3 |
| Tez       | org.apache.tez       | tez-javadoc-tools                   | 0.9.1.7.1.9.6-3 |
| Tez       | org.apache.tez       | tez-job-analyzer                    | 0.9.1.7.1.9.6-3 |
| Tez       | org.apache.tez       | tez-mapreduce                       | 0.9.1.7.1.9.6-3 |
| Tez       | org.apache.tez       | tez-protobuf-history-plugin         | 0.9.1.7.1.9.6-3 |
| Tez       | org.apache.tez       | tez-runtime-internals               | 0.9.1.7.1.9.6-3 |
| Tez       | org.apache.tez       | tez-runtime-library                 | 0.9.1.7.1.9.6-3 |
| Tez       | org.apache.tez       | tez-tests                           | 0.9.1.7.1.9.6-3 |
| Tez       | org.apache.tez       | tez-yarn-timeline-cache-plugin      | 0.9.1.7.1.9.6-3 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history           | 0.9.1.7.1.9.6-3 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history-with-acls | 0.9.1.7.1.9.6-3 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history-with-fs   | 0.9.1.7.1.9.6-3 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-angular                    | 0.8.2.7.1.9.6-3 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-display                    | 0.8.2.7.1.9.6-3 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-interpreter                | 0.8.2.7.1.9.6-3 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-jdbc                       | 0.8.2.7.1.9.6-3 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-jupyter                    | 0.8.2.7.1.9.6-3 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-livy                       | 0.8.2.7.1.9.6-3 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-markdown                   | 0.8.2.7.1.9.6-3 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-server                     | 0.8.2.7.1.9.6-3 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-shell                      | 0.8.2.7.1.9.6-3 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-zengine                    | 0.8.2.7.1.9.6-3 |
| ZooKeeper | org.apache.zookeeper | zookeeper                           | 3.8.1.7.1.9.6-3 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-fatjar            | 3.8.1.7.1.9.6-3 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-loggraph          | 3.8.1.7.1.9.6-3 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-rest              | 3.8.1.7.1.9.6-3 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-zooinspector      | 3.8.1.7.1.9.6-3 |
| ZooKeeper | org.apache.zookeeper | zookeeper-it                        | 3.8.1.7.1.9.6-3 |
| ZooKeeper | org.apache.zookeeper | zookeeper-jute                      | 3.8.1.7.1.9.6-3 |
| ZooKeeper | org.apache.zookeeper | zookeeper-prometheus-metrics        | 3.8.1.7.1.9.6-3 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-election          | 3.8.1.7.1.9.6-3 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-lock              | 3.8.1.7.1.9.6-3 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-queue             | 3.8.1.7.1.9.6-3 |

## Runtime 7.1.9.7-2

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 | 3.0.0.7.1.9.7-2 |
| Atlas   | org.apache.atlas | atlas-aws-s3-bridge | 3.0.0.7.1.9.7-2 |

| Project | groupId          | artifactId                   | version          |
|---------|------------------|------------------------------|------------------|
| Atlas   | org.apache.atlas | atlas-azure-adls-bridge      | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-classification-updater | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-client-common          | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-client-v1              | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-client-v2              | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-common                 | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-distro                 | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-docs                   | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-graphdb-api            | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-graphdb-common         | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-graphdb-janus          | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-hdfs-bridge            | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-index-repair-tool      | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-intg                   | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-janusgraph-hbase2      | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-notification           | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-plugin-classloader     | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-repository             | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-server-api             | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | atlas-testtools              | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | hbase-bridge                 | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | hbase-bridge-shim            | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | hbase-testing-util           | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | hdfs-model                   | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | hive-bridge                  | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | hive-bridge-shim             | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | impala-bridge                | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | impala-bridge-shim           | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | impala-hook-api              | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | kafka-bridge                 | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | kafka-bridge-shim            | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | navigator-to-atlas           | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | sample-app                   | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | sqoop-bridge                 | 3.0.0.7.1.9.7-2  |
| Atlas   | org.apache.atlas | sqoop-bridge-shim            | 3.0.0.7.1.9.7-2  |
| Avro    | org.apache.avro  | avro                         | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro  | avro-android                 | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro  | avro-codegen-test            | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro  | avro-compiler                | 1.11.1.7.1.9.7-2 |

| Project | groupId            | artifactId                      | version          |
|---------|--------------------|---------------------------------|------------------|
| Avro    | org.apache.avro    | avro-grpc                       | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro    | avro-ipc                        | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro    | avro-ipc-jetty                  | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro    | avro-ipc-netty                  | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro    | avro-mapred                     | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro    | avro-maven-plugin               | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro    | avro-perf                       | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro    | avro-protobuf                   | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro    | avro-service-archetype          | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro    | avro-test-custom-conversions    | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro    | avro-thrift                     | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro    | avro-tools                      | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro    | trevni-avro                     | 1.11.1.7.1.9.7-2 |
| Avro    | org.apache.avro    | trevni-core                     | 1.11.1.7.1.9.7-2 |
| Calcite | org.apache.calcite | calcite-babel                   | 1.19.0.7.1.9.7-2 |
| Calcite | org.apache.calcite | calcite-core                    | 1.19.0.7.1.9.7-2 |
| Calcite | org.apache.calcite | calcite-druid                   | 1.19.0.7.1.9.7-2 |
| Calcite | org.apache.calcite | calcite-linq4j                  | 1.19.0.7.1.9.7-2 |
| Calcite | org.apache.calcite | calcite-server                  | 1.19.0.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop  | hadoop-aliyun                   | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-annotations              | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-archive-logs             | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-archives                 | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-assemblies               | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-auth                     | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-aws                      | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-azure                    | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-azure-datalake           | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-build-tools              | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-client                   | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-client-api               | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-client-integration-tests | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-client-minicluster       | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-client-runtime           | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-cloud-storage            | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-common                   | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-datajoin                 | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-distcp                   | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop  | hadoop-extras                   | 3.1.1.7.1.9.7-2  |

| Project | groupId           | artifactId                                     | version         |
|---------|-------------------|------------------------------------------------|-----------------|
| Hadoop  | org.apache.hadoop | hadoop-fs2img                                  | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-gridmix                                 | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs                                    | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-client                             | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-httpfs                             | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-native-client                      | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-nfs                                | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-rbf                                | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-kafka                                   | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-kms                                     | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-app                    | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-common                 | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-core                   | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-hs                     | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-hs-plugins             | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-jobclient              | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-nativetask             | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-shuffle                | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-uploader               | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-examples                      | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-maven-plugins                           | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-minicluster                             | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-minikdc                                 | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-nfs                                     | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-openstack                               | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-resourceestimator                       | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-rumen                                   | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-sls                                     | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-streaming                               | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-tools-dist                              | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-api                                | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-distributedshell      | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-unmanaged-am-launcher | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-client                             | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-common                             | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-registry                           | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-applicationhistoryservice   | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-common                      | 3.1.1.7.1.9.7-2 |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-nodemanager                 | 3.1.1.7.1.9.7-2 |

| Project | groupId           | artifactId                                        | version          |
|---------|-------------------|---------------------------------------------------|------------------|
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-resourcemanager                | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-router                         | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-sharedcachemanager             | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-tests                          | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timeline-pluginstorage         | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice                | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-client   | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-common   | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-server-2 | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-tests    | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-web-proxy                      | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-api                          | 3.1.1.7.1.9.7-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-core                         | 3.1.1.7.1.9.7-2  |
| HBase   | org.apache.hbase  | hbase-annotations                                 | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-asyncfs                                     | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-checkstyle                                  | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-client                                      | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-client-project                              | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-common                                      | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-endpoint                                    | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-examples                                    | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-external-blockcache                         | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-hadoop-compat                               | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-hadoop2-compat                              | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-hbtop                                       | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-http                                        | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-it                                          | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-logging                                     | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-mapreduce                                   | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-metrics                                     | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-metrics-api                                 | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-procedure                                   | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-protocol                                    | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-protocol-shaded                             | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-replication                                 | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-resource-bundle                             | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-rest                                        | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-rsgroup                                     | 2.4.17.7.1.9.7-2 |
| HBase   | org.apache.hbase  | hbase-server                                      | 2.4.17.7.1.9.7-2 |

| Project | groupId          | artifactId                       | version            |
|---------|------------------|----------------------------------|--------------------|
| HBase   | org.apache.hbase | hbase-shaded-client              | 2.4.17.7.1.9.7-2   |
| HBase   | org.apache.hbase | hbase-shaded-client-byo-hadoop   | 2.4.17.7.1.9.7-2   |
| HBase   | org.apache.hbase | hbase-shaded-client-project      | 2.4.17.7.1.9.7-2   |
| HBase   | org.apache.hbase | hbase-shaded-mapreduce           | 2.4.17.7.1.9.7-2   |
| HBase   | org.apache.hbase | hbase-shaded-testing-util        | 2.4.17.7.1.9.7-2   |
| HBase   | org.apache.hbase | hbase-shaded-testing-util-tester | 2.4.17.7.1.9.7-2   |
| HBase   | org.apache.hbase | hbase-shell                      | 2.4.17.7.1.9.7-2   |
| HBase   | org.apache.hbase | hbase-testing-util               | 2.4.17.7.1.9.7-2   |
| HBase   | org.apache.hbase | hbase-thrift                     | 2.4.17.7.1.9.7-2   |
| HBase   | org.apache.hbase | hbase-zookeeper                  | 2.4.17.7.1.9.7-2   |
| Hive    | org.apache.hive  | hive-beeline                     | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-blobstore                   | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-classification              | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-cli                         | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-common                      | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-contrib                     | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-exec                        | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-hbase-handler               | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-hcatalog-it-unit            | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-hplsql                      | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-it-custom-serde             | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-it-minikdc                  | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-it-qfile                    | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-it-qfile-kudu               | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-it-test-serde               | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-it-unit                     | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-it-unit-hadoop2             | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-it-util                     | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-jdbc                        | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-jdbc-handler                | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-jmh                         | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-kryo-registrator            | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-kudu-handler                | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-llap-client                 | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-llap-common                 | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-llap-ext-client             | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-llap-server                 | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-llap-tez                    | 3.1.3000.7.1.9.7-2 |
| Hive    | org.apache.hive  | hive-metastore                   | 3.1.3000.7.1.9.7-2 |

| Project                  | groupId              | artifactId                               | version            |
|--------------------------|----------------------|------------------------------------------|--------------------|
| Hive                     | org.apache.hive      | hive-pre-upgrade                         | 3.1.3000.7.1.9.7-2 |
| Hive                     | org.apache.hive      | hive-serde                               | 3.1.3000.7.1.9.7-2 |
| Hive                     | org.apache.hive      | hive-service                             | 3.1.3000.7.1.9.7-2 |
| Hive                     | org.apache.hive      | hive-service-rpc                         | 3.1.3000.7.1.9.7-2 |
| Hive                     | org.apache.hive      | hive-shims                               | 3.1.3000.7.1.9.7-2 |
| Hive                     | org.apache.hive      | hive-spark-client                        | 3.1.3000.7.1.9.7-2 |
| Hive                     | org.apache.hive      | hive-standalone-metastore                | 3.1.3000.7.1.9.7-2 |
| Hive                     | org.apache.hive      | hive-storage-api                         | 3.1.3000.7.1.9.7-2 |
| Hive                     | org.apache.hive      | hive-streaming                           | 3.1.3000.7.1.9.7-2 |
| Hive                     | org.apache.hive      | hive-testutils                           | 3.1.3000.7.1.9.7-2 |
| Hive                     | org.apache.hive      | hive-vector-code-gen                     | 3.1.3000.7.1.9.7-2 |
| Hive                     | org.apache.hive      | kafka-handler                            | 3.1.3000.7.1.9.7-2 |
| Hive Warehouse Connector | com.hortonworks.hive | hive-warehouse-connector_2.11            | 1.0.0.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | ci                                       | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | connect                                  | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | connect-api                              | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | connect-basic-auth-extension             | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | connect-cloudera-authorization-extension | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | connect-cloudera-common                  | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | connect-cloudera-secret-storage          | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | connect-cloudera-security-policies       | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | connect-file                             | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | connect-json                             | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | connect-mirror                           | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | connect-mirror-client                    | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | connect-runtime                          | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | connect-transforms                       | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | generator                                | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | kafka-clients                            | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | kafka-cloudera-metrics-reporter_2.12     | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | kafka-cloudera-metrics-reporter_2.13     | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | kafka-cloudera-plugins                   | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | kafka-examples                           | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | kafka-group-coordinator                  | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | kafka-log4j-appender                     | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | kafka-metadata                           | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | kafka-raft                               | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | kafka-server-common                      | 3.4.1.7.1.9.7-2    |
| Kafka                    | org.apache.kafka     | kafka-shell                              | 3.4.1.7.1.9.7-2    |



| Project | groupId          | artifactId                              | version         |
|---------|------------------|-----------------------------------------|-----------------|
| Kafka   | org.apache.kafka | kafka-storage                           | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-storage-api                       | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams                           | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-examples                  | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-scala_2.12                | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-scala_2.13                | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-test-utils                | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0100 | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0101 | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0102 | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0110 | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-10   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-11   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-20   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-21   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-22   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-23   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-24   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-25   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-26   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-27   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-28   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-30   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-31   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-32   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-33   | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka-tools                             | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka_2.12                              | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | kafka_2.13                              | 3.4.1.7.1.9.7-2 |
| Kafka   | org.apache.kafka | trogdor                                 | 3.4.1.7.1.9.7-2 |
| Knox    | org.apache.knox  | gateway-adapter                         | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox  | gateway-admin-ui                        | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox  | gateway-applications                    | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox  | gateway-cloud-bindings                  | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox  | gateway-demo-ldap                       | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox  | gateway-demo-ldap-launcher              | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox  | gateway-discovery-ambari                | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox  | gateway-discovery-cm                    | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox  | gateway-docker                          | 1.3.0.7.1.9.7-2 |

| Project | groupId         | artifactId                                        | version         |
|---------|-----------------|---------------------------------------------------|-----------------|
| Knox    | org.apache.knox | gateway-i18n                                      | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-i18n-logging-log4j                        | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-i18n-logging-sl4j                         | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-performance-test                          | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-ha                               | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-common        | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-concat        | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-hadoop-groups | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-pseudo        | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-regex         | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-switchcase    | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-jersey                           | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-rewrite                          | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-rewrite-common                   | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-hostmap-static      | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-inbound-query-param | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-service-registry    | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-rewrite-step-encrypt-uri         | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-rewrite-step-secure-query        | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-security-authc-anon              | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-security-authz-acls              | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-security-authz-composite         | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-security-clientcert              | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-security-hadoopauth              | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-security-jwt                     | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-security-pac4j                   | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-security-preauth                 | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-security-shiro                   | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-provider-security-webappsec               | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-release                                   | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-server                                    | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-server-launcher                           | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-server-xforwarded-filter                  | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-admin                             | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-as                                | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-definitions                       | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-hashicorp-vault                   | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-hbase                             | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-health                            | 1.3.0.7.1.9.7-2 |

| Project | groupId         | artifactId                    | version         |
|---------|-----------------|-------------------------------|-----------------|
| Knox    | org.apache.knox | gateway-service-hive          | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-idbroker      | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-impala        | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-jkg           | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-knoxsso       | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-knoxssout     | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-knoxtoken     | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-livy          | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-metadata      | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-nifi          | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-nifi-registry | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-remoteconfig  | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-rm            | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-session       | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-storm         | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-test          | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-tgs           | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-vault         | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-service-webhdfs       | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-shell                 | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-shell-launcher        | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-shell-release         | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-shell-samples         | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-spi                   | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-test                  | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-test-idbroker         | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-test-release-utils    | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-test-utils            | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-topology-hadoop-xml   | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-topology-simple       | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-util-common           | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-util-configinjector   | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-util-launcher         | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | gateway-util-urltemplate      | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | hadoop-examples               | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | knox-cli-launcher             | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | knox-homepage-ui              | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | knox-token-generation-ui      | 1.3.0.7.1.9.7-2 |
| Knox    | org.apache.knox | knox-token-management-ui      | 1.3.0.7.1.9.7-2 |

| Project | groupId           | artifactId                  | version          |
|---------|-------------------|-----------------------------|------------------|
| Knox    | org.apache.knox   | webhdfs-kerb-test           | 1.3.0.7.1.9.7-2  |
| Knox    | org.apache.knox   | webhdfs-test                | 1.3.0.7.1.9.7-2  |
| Kudu    | org.apache.kudu   | kudu-backup-tools           | 1.17.0.7.1.9.7-2 |
| Kudu    | org.apache.kudu   | kudu-backup2_2.11           | 1.17.0.7.1.9.7-2 |
| Kudu    | org.apache.kudu   | kudu-backup3_2.12           | 1.17.0.7.1.9.7-2 |
| Kudu    | org.apache.kudu   | kudu-client                 | 1.17.0.7.1.9.7-2 |
| Kudu    | org.apache.kudu   | kudu-hive                   | 1.17.0.7.1.9.7-2 |
| Kudu    | org.apache.kudu   | kudu-spark2-tools_2.11      | 1.17.0.7.1.9.7-2 |
| Kudu    | org.apache.kudu   | kudu-spark2_2.11            | 1.17.0.7.1.9.7-2 |
| Kudu    | org.apache.kudu   | kudu-spark3-tools_2.12      | 1.17.0.7.1.9.7-2 |
| Kudu    | org.apache.kudu   | kudu-spark3_2.12            | 1.17.0.7.1.9.7-2 |
| Kudu    | org.apache.kudu   | kudu-test-utils             | 1.17.0.7.1.9.7-2 |
| Livy    | org.apache.livy   | livy-api                    | 0.7.2.7.1.9.7-2  |
| Livy    | org.apache.livy   | livy-client-common          | 0.7.2.7.1.9.7-2  |
| Livy    | org.apache.livy   | livy-client-http            | 0.7.2.7.1.9.7-2  |
| Livy    | org.apache.livy   | livy-core_2.11              | 0.7.2.7.1.9.7-2  |
| Livy    | org.apache.livy   | livy-examples               | 0.7.2.7.1.9.7-2  |
| Livy    | org.apache.livy   | livy-integration-test       | 0.7.2.7.1.9.7-2  |
| Livy    | org.apache.livy   | livy-repl_2.11              | 0.7.2.7.1.9.7-2  |
| Livy    | org.apache.livy   | livy-rsc                    | 0.7.2.7.1.9.7-2  |
| Livy    | org.apache.livy   | livy-scala-api_2.11         | 0.7.2.7.1.9.7-2  |
| Livy    | org.apache.livy   | livy-server                 | 0.7.2.7.1.9.7-2  |
| Livy    | org.apache.livy   | livy-test-lib               | 0.7.2.7.1.9.7-2  |
| Livy    | org.apache.livy   | livy-thriftserver           | 0.7.2.7.1.9.7-2  |
| Livy    | org.apache.livy   | livy-thriftserver-session   | 0.7.2.7.1.9.7-2  |
| Lucene  | org.apache.lucene | lucene-analyzers-common     | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-icu        | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-kuromoji   | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-morfologik | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-nori       | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-openssl    | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-phonetic   | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-smartcn    | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-stempel    | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-backward-codecs      | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-benchmark            | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-classification       | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-codecs               | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-core                 | 8.11.2.7.1.9.7-2 |

| Project | groupId           | artifactId                     | version          |
|---------|-------------------|--------------------------------|------------------|
| Lucene  | org.apache.lucene | lucene-demo                    | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-expressions             | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-facet                   | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-grouping                | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-highlighter             | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-join                    | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-memory                  | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-misc                    | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-monitor                 | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-queries                 | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-queryparser             | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-replicator              | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-sandbox                 | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-spatial-extras          | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-spatial3d               | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-suggest                 | 8.11.2.7.1.9.7-2 |
| Lucene  | org.apache.lucene | lucene-test-framework          | 8.11.2.7.1.9.7-2 |
| Oozie   | org.apache.oozie  | oozie-client                   | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-core                     | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-distro                   | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-examples                 | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-api           | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-client        | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-server                   | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-sharelib-distcp          | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-sharelib-git             | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hcatalog        | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hive            | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hive2           | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-sharelib-oozie           | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-sharelib-spark           | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-sharelib-spark3          | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-sharelib-sqoop           | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-sharelib-streaming       | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-tools                    | 5.1.0.7.1.9.7-2  |
| Oozie   | org.apache.oozie  | oozie-zookeeper-security-tests | 5.1.0.7.1.9.7-2  |
| ORC     | org.apache.orc    | orc-core                       | 1.5.1.7.1.9.7-2  |
| ORC     | org.apache.orc    | orc-examples                   | 1.5.1.7.1.9.7-2  |
| ORC     | org.apache.orc    | orc-mapreduce                  | 1.5.1.7.1.9.7-2  |

| Project | groupId          | artifactId                    | version         |
|---------|------------------|-------------------------------|-----------------|
| ORC     | org.apache.orc   | orc-shims                     | 1.5.1.7.1.9.7-2 |
| ORC     | org.apache.orc   | orc-tools                     | 1.5.1.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-annotation-processing    | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-client                   | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-common                   | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-config                   | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-container-service        | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-docs                     | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-erasurecode              | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-client | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-server | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-test   | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-interface-admin          | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-interface-client         | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-interface-server         | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-managed-rocksdb          | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-rocks-native             | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-server-framework         | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-server-scm               | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-test-utils               | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | hdds-tools                    | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | mini-chaos-tests              | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-client                  | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-common                  | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-csi                     | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-datanode                | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-dist                    | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-filesystem              | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-filesystem-common       | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-filesystem-hadoop2      | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-filesystem-hadoop3      | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-filesystem-shaded       | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-httpsgateway            | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-insight                 | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-integration-test        | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-interface-client        | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-interface-storage       | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-manager                 | 1.4.0.7.1.9.7-2 |
| Ozone   | org.apache.ozone | ozone-network-tests           | 1.4.0.7.1.9.7-2 |

| Project | groupId            | artifactId                             | version           |
|---------|--------------------|----------------------------------------|-------------------|
| Ozone   | org.apache.ozone   | ozone-recon                            | 1.4.0.7.1.9.7-2   |
| Ozone   | org.apache.ozone   | ozone-reconcodegen                     | 1.4.0.7.1.9.7-2   |
| Ozone   | org.apache.ozone   | ozone-s3-secret-store                  | 1.4.0.7.1.9.7-2   |
| Ozone   | org.apache.ozone   | ozone-s3gateway                        | 1.4.0.7.1.9.7-2   |
| Ozone   | org.apache.ozone   | ozone-tools                            | 1.4.0.7.1.9.7-2   |
| Ozone   | org.apache.ozone   | rocksdb-checkpoint-differ              | 1.4.0.7.1.9.7-2   |
| Parquet | org.apache.parquet | parquet-avro                           | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-cascading                      | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-cascading3                     | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-column                         | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-common                         | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-encoding                       | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-format-structures              | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-generator                      | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-hadoop                         | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-hadoop-bundle                  | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-jackson                        | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-pig                            | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-pig-bundle                     | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-protobuf                       | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-scala_2.10                     | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-thrift                         | 1.10.99.7.1.9.7-2 |
| Parquet | org.apache.parquet | parquet-tools                          | 1.10.99.7.1.9.7-2 |
| Phoenix | org.apache.phoenix | phoenix-client-embedded-hbase-2.4      | 5.1.1.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-client-hbase-2.4               | 5.1.1.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-connectors-phoenix5-compatible | 6.0.0.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-core                           | 5.1.1.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.1.6         | 5.1.1.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.2.5         | 5.1.1.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.3.0         | 5.1.1.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.4.0         | 5.1.1.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compatible-2.4.1         | 5.1.1.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-pherf                          | 5.1.1.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-queryserver                    | 6.0.0.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-client             | 6.0.0.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-it                 | 6.0.0.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-load-balancer      | 6.0.0.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-queryserver-orchestrator       | 6.0.0.7.1.9.7-2   |
| Phoenix | org.apache.phoenix | phoenix-server-hbase-2.4               | 5.1.1.7.1.9.7-2   |

| Project | groupId            | artifactId                  | version         |
|---------|--------------------|-----------------------------|-----------------|
| Phoenix | org.apache.phoenix | phoenix-tracing-webapp      | 5.1.1.7.1.9.7-2 |
| Phoenix | org.apache.phoenix | phoenix5-hive               | 6.0.0.7.1.9.7-2 |
| Phoenix | org.apache.phoenix | phoenix5-hive-shaded        | 6.0.0.7.1.9.7-2 |
| Phoenix | org.apache.phoenix | phoenix5-spark              | 6.0.0.7.1.9.7-2 |
| Phoenix | org.apache.phoenix | phoenix5-spark-shaded       | 6.0.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | conditions-enrichers        | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | credentialbuilder           | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | embeddedwebserver           | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | jisql                       | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ldapconfigcheck             | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-adls-plugin          | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-atlas-plugin         | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-atlas-plugin-shim    | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-authn                | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-common-ha            | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-distro               | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-examples-distro      | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-hbase-plugin         | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-hbase-plugin-shim    | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-hdfs-plugin          | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-hdfs-plugin-shim     | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-hive-plugin          | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-hive-plugin-shim     | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-intg                 | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-kafka-connect-plugin | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-kafka-plugin         | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-kafka-plugin-shim    | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-kms                  | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-kms-plugin           | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-kms-plugin-shim      | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-knox-plugin          | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-knox-plugin-shim     | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-kudu-plugin          | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-kylin-plugin         | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-kylin-plugin-shim    | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-metrics              | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-nifi-plugin          | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-nifi-registry-plugin | 2.4.0.7.1.9.7-2 |
| Ranger  | org.apache.ranger  | ranger-ozone-plugin         | 2.4.0.7.1.9.7-2 |



| Project | groupId           | artifactId                    | version          |
|---------|-------------------|-------------------------------|------------------|
| Ranger  | org.apache.ranger | ranger-ozone-plugin-shim      | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-plugin-classloader     | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-plugins-audit          | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-plugins-common         | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-plugins-cred           | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-plugins-installer      | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-policymigration        | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-raz-adls               | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-raz-chained-plugins    | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-raz-hook-abfs          | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-raz-intg               | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-raz-processor          | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-rms-common             | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-rms-hive               | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-rms-plugins-common     | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-rms-tools              | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-rms-webapp             | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-sampleapp-plugin       | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-schema-registry-plugin | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-solr-plugin            | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-solr-plugin-shim       | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin           | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin-shim      | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-storm-plugin           | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-storm-plugin-shim      | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-tagsync                | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-tools                  | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-util                   | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-yarn-plugin            | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ranger-yarn-plugin-shim       | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | sample-client                 | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | sampleapp                     | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | shaded-raz-hook-abfs          | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | ugsync-util                   | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | unixauthclient                | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | unixauthservice               | 2.4.0.7.1.9.7-2  |
| Ranger  | org.apache.ranger | unixusersync                  | 2.4.0.7.1.9.7-2  |
| Solr    | org.apache.solr   | solr-analysis-extras          | 8.11.2.7.1.9.7-2 |
| Solr    | org.apache.solr   | solr-analytics                | 8.11.2.7.1.9.7-2 |

| Project | groupId          | artifactId                               | version          |
|---------|------------------|------------------------------------------|------------------|
| Solr    | org.apache.solr  | solr-cell                                | 8.11.2.7.1.9.7-2 |
| Solr    | org.apache.solr  | solr-core                                | 8.11.2.7.1.9.7-2 |
| Solr    | org.apache.solr  | solr-dataimporthandler                   | 8.11.2.7.1.9.7-2 |
| Solr    | org.apache.solr  | solr-dataimporthandler-extras            | 8.11.2.7.1.9.7-2 |
| Solr    | org.apache.solr  | solr-gcs-repository                      | 8.11.2.7.1.9.7-2 |
| Solr    | org.apache.solr  | solr-jaegertracer-configurator           | 8.11.2.7.1.9.7-2 |
| Solr    | org.apache.solr  | solr-langid                              | 8.11.2.7.1.9.7-2 |
| Solr    | org.apache.solr  | solr-ltr                                 | 8.11.2.7.1.9.7-2 |
| Solr    | org.apache.solr  | solr-prometheus-exporter                 | 8.11.2.7.1.9.7-2 |
| Solr    | org.apache.solr  | solr-s3-repository                       | 8.11.2.7.1.9.7-2 |
| Solr    | org.apache.solr  | solr-security-util                       | 8.11.2.7.1.9.7-2 |
| Solr    | org.apache.solr  | solr-solrj                               | 8.11.2.7.1.9.7-2 |
| Solr    | org.apache.solr  | solr-test-framework                      | 8.11.2.7.1.9.7-2 |
| Solr    | org.apache.solr  | solr-velocity                            | 8.11.2.7.1.9.7-2 |
| Spark   | org.apache.spark | spark-avro_2.11                          | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-catalyst_2.11                      | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-core_2.11                          | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-graphx_2.11                        | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-hadoop-cloud_2.11                  | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-hive_2.11                          | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-kubernetes_2.11                    | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-kvstore_2.11                       | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-launcher_2.11                      | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-mllib-local_2.11                   | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-mllib_2.11                         | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-network-common_2.11                | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-network-shuffle_2.11               | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-network-yarn_2.11                  | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-repl_2.11                          | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-sketch_2.11                        | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-sql-kafka-0-10_2.11                | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-sql_2.11                           | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-streaming-kafka-0-10-assembly_2.11 | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-streaming-kafka-0-10_2.11          | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-streaming_2.11                     | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-tags_2.11                          | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-token-provider-kafka-0-10_2.11     | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-unsafe_2.11                        | 2.4.8.7.1.9.7-2  |
| Spark   | org.apache.spark | spark-yarn_2.11                          | 2.4.8.7.1.9.7-2  |

| Project   | groupId              | artifactId                          | version         |
|-----------|----------------------|-------------------------------------|-----------------|
| Sqoop     | org.apache.sqoop     | sqoop                               | 1.4.7.7.1.9.7-2 |
| Sqoop     | org.apache.sqoop     | sqoop-test                          | 1.4.7.7.1.9.7-2 |
| Tez       | org.apache.tez       | hadoop-shim                         | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | hadoop-shim-2.8                     | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-api                             | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-aux-services                    | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-common                          | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-dag                             | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-examples                        | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-ext-service-tests               | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-history-parser                  | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-javadoc-tools                   | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-job-analyzer                    | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-mapreduce                       | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-protobuf-history-plugin         | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-runtime-internals               | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-runtime-library                 | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-tests                           | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-yarn-timeline-cache-plugin      | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history           | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history-with-acls | 0.9.1.7.1.9.7-2 |
| Tez       | org.apache.tez       | tez-yarn-timeline-history-with-fs   | 0.9.1.7.1.9.7-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-angular                    | 0.8.2.7.1.9.7-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-display                    | 0.8.2.7.1.9.7-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-interpreter                | 0.8.2.7.1.9.7-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-jdbc                       | 0.8.2.7.1.9.7-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-jupyter                    | 0.8.2.7.1.9.7-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-livy                       | 0.8.2.7.1.9.7-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-markdown                   | 0.8.2.7.1.9.7-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-server                     | 0.8.2.7.1.9.7-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-shell                      | 0.8.2.7.1.9.7-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-zengine                    | 0.8.2.7.1.9.7-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper                           | 3.8.1.7.1.9.7-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-fatjar            | 3.8.1.7.1.9.7-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-loggraph          | 3.8.1.7.1.9.7-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-rest              | 3.8.1.7.1.9.7-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-zooinspector      | 3.8.1.7.1.9.7-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-it                        | 3.8.1.7.1.9.7-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-jute                      | 3.8.1.7.1.9.7-2 |

| Project   | groupId              | artifactId                   | version         |
|-----------|----------------------|------------------------------|-----------------|
| ZooKeeper | org.apache.zookeeper | zookeeper-prometheus-metrics | 3.8.1.7.1.9.7-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-election   | 3.8.1.7.1.9.7-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-lock       | 3.8.1.7.1.9.7-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-queue      | 3.8.1.7.1.9.7-2 |

### Runtime 7.1.9.9-1

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          | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-aws-s3-bridge          | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-azure-adls-bridge      | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-classification-updater | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-client-common          | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-client-v1              | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-client-v2              | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-common                 | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-distro                 | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-docs                   | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-graphdb-api            | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-graphdb-common         | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-graphdb-janus          | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-hdfs-bridge            | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-index-repair-tool      | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-intg                   | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-janusgraph-hbase2      | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-notification           | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-plugin-classloader     | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-repository             | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-server-api             | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | atlas-testtools              | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | hbase-bridge                 | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | hbase-bridge-shim            | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | hbase-testing-util           | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | hdfs-model                   | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | hive-bridge                  | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | hive-bridge-shim             | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | impala-bridge                | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | impala-bridge-shim           | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | impala-hook-api              | 3.0.0.7.1.9.9-1 |
| Atlas   | org.apache.atlas | kafka-bridge                 | 3.0.0.7.1.9.9-1 |

| Project | groupId            | artifactId                   | version          |
|---------|--------------------|------------------------------|------------------|
| Atlas   | org.apache.atlas   | kafka-bridge-shim            | 3.0.0.7.1.9.9-1  |
| Atlas   | org.apache.atlas   | navigator-to-atlas           | 3.0.0.7.1.9.9-1  |
| Atlas   | org.apache.atlas   | sample-app                   | 3.0.0.7.1.9.9-1  |
| Atlas   | org.apache.atlas   | sqoop-bridge                 | 3.0.0.7.1.9.9-1  |
| Atlas   | org.apache.atlas   | sqoop-bridge-shim            | 3.0.0.7.1.9.9-1  |
| Avro    | org.apache.avro    | avro                         | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-android                 | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-codegen-test            | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-compiler                | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-grpc                    | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-ipc                     | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-ipc-jetty               | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-ipc-netty               | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-mapred                  | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-maven-plugin            | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-perf                    | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-protobuf                | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-service-archetype       | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-test-custom-conversions | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-thrift                  | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | avro-tools                   | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | trevni-avro                  | 1.11.1.7.1.9.9-1 |
| Avro    | org.apache.avro    | trevni-core                  | 1.11.1.7.1.9.9-1 |
| Calcite | org.apache.calcite | calcite-babel                | 1.19.0.7.1.9.9-1 |
| Calcite | org.apache.calcite | calcite-core                 | 1.19.0.7.1.9.9-1 |
| Calcite | org.apache.calcite | calcite-druid                | 1.19.0.7.1.9.9-1 |
| Calcite | org.apache.calcite | calcite-linq4j               | 1.19.0.7.1.9.9-1 |
| Calcite | org.apache.calcite | calcite-server               | 1.19.0.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop  | hadoop-aliyun                | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop  | hadoop-annotations           | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop  | hadoop-archive-logs          | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop  | hadoop-archives              | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop  | hadoop-assemblies            | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop  | hadoop-auth                  | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop  | hadoop-aws                   | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop  | hadoop-azure                 | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop  | hadoop-azure-datalake        | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop  | hadoop-build-tools           | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop  | hadoop-client                | 3.1.1.7.1.9.9-1  |

| Project | groupId           | artifactId                         | version         |
|---------|-------------------|------------------------------------|-----------------|
| Hadoop  | org.apache.hadoop | hadoop-client-api                  | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-client-integration-tests    | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-client-minicluster          | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-client-runtime              | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-cloud-storage               | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-common                      | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-datajoin                    | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-distcp                      | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-extras                      | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-fs2img                      | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-gridmix                     | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs                        | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-client                 | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-httpfs                 | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-native-client          | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-nfs                    | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-rbf                    | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-kafka                       | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-kms                         | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-app        | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-common     | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-core       | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-hs         | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-hs-plugins | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-jobclient  | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-nativetask | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-shuffle    | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-uploader   | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-examples          | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-maven-plugins               | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-minicluster                 | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-minikdc                     | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-nfs                         | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-openstack                   | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-resourceestimator           | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-rumen                       | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-sls                         | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-streaming                   | 3.1.1.7.1.9.9-1 |
| Hadoop  | org.apache.hadoop | hadoop-tools-dist                  | 3.1.1.7.1.9.9-1 |

| Project | groupId           | artifactId                                        | version          |
|---------|-------------------|---------------------------------------------------|------------------|
| Hadoop  | org.apache.hadoop | hadoop-yarn-api                                   | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-distributedshell         | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-unmanaged-am-launcher    | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-client                                | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-common                                | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-registry                              | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-applicationhistoryservice      | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-common                         | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-nodemanager                    | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-resourcemanager                | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-router                         | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-sharedcachemanager             | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-tests                          | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timeline-pluginstorage         | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice                | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-client   | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-common   | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-server-2 | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-tests    | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-web-proxy                      | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-api                          | 3.1.1.7.1.9.9-1  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-core                         | 3.1.1.7.1.9.9-1  |
| HBase   | org.apache.hbase  | hbase-annotations                                 | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-asyncfs                                     | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-checkstyle                                  | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-client                                      | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-client-project                              | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-common                                      | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-endpoint                                    | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-examples                                    | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-external-blockcache                         | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-hadoop-compat                               | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-hadoop2-compat                              | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-hbtop                                       | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-http                                        | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-it                                          | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-logging                                     | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-mapreduce                                   | 2.4.17.7.1.9.9-1 |
| HBase   | org.apache.hbase  | hbase-metrics                                     | 2.4.17.7.1.9.9-1 |

| Project | groupId          | artifactId                       | version            |
|---------|------------------|----------------------------------|--------------------|
| HBase   | org.apache.hbase | hbase-metrics-api                | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-procedure                  | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-protocol                   | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-protocol-shaded            | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-replication                | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-resource-bundle            | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-rest                       | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-rsgroup                    | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-server                     | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-shaded-client              | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-shaded-client-byo-hadoop   | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-shaded-client-project      | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-shaded-mapreduce           | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-shaded-testing-util        | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-shaded-testing-util-tester | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-shell                      | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-testing-util               | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-thrift                     | 2.4.17.7.1.9.9-1   |
| HBase   | org.apache.hbase | hbase-zookeeper                  | 2.4.17.7.1.9.9-1   |
| Hive    | org.apache.hive  | hive-beeline                     | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-blobstore                   | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-classification              | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-cli                         | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-common                      | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-contrib                     | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-exec                        | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-hbase-handler               | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-hcatalog-it-unit            | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-hplsql                      | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-it-custom-serde             | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-it-minikdc                  | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-it-qfile                    | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-it-qfile-kudu               | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-it-test-serde               | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-it-unit                     | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-it-unit-hadoop2             | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-it-util                     | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-jdbc                        | 3.1.3000.7.1.9.9-1 |
| Hive    | org.apache.hive  | hive-jdbc-handler                | 3.1.3000.7.1.9.9-1 |



| Project                  | groupId              | artifactId                               | version            |
|--------------------------|----------------------|------------------------------------------|--------------------|
| Hive                     | org.apache.hive      | hive-jmh                                 | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-kryo-registrator                    | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-kudu-handler                        | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-llap-client                         | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-llap-common                         | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-llap-ext-client                     | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-llap-server                         | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-llap-tez                            | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-metastore                           | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-pre-upgrade                         | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-serde                               | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-service                             | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-service-rpc                         | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-shims                               | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-spark-client                        | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-standalone-metastore                | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-storage-api                         | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-streaming                           | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-testutils                           | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | hive-vector-code-gen                     | 3.1.3000.7.1.9.9-1 |
| Hive                     | org.apache.hive      | kafka-handler                            | 3.1.3000.7.1.9.9-1 |
| Hive Warehouse Connector | com.hortonworks.hive | hive-warehouse-connector_2.11            | 1.0.0.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | ci                                       | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | connect                                  | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | connect-api                              | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | connect-basic-auth-extension             | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | connect-cloudera-authorization-extension | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | connect-cloudera-common                  | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | connect-cloudera-secret-storage          | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | connect-cloudera-security-policies       | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | connect-file                             | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | connect-json                             | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | connect-mirror                           | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | connect-mirror-client                    | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | connect-runtime                          | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | connect-transforms                       | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | generator                                | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | kafka-clients                            | 3.4.1.7.1.9.9-1    |
| Kafka                    | org.apache.kafka     | kafka-cloudera-metrics-reporter_2.12     | 3.4.1.7.1.9.9-1    |

| Project | groupId          | artifactId                              | version         |
|---------|------------------|-----------------------------------------|-----------------|
| Kafka   | org.apache.kafka | kafka-cloudera-metrics-reporter_2.13    | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-cloudera-plugins                  | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-examples                          | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-group-coordinator                 | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-log4j-appender                    | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-metadata                          | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-raft                              | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-server-common                     | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-shell                             | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-storage                           | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-storage-api                       | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams                           | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-examples                  | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-scala_2.12                | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-scala_2.13                | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-test-utils                | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0100 | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0101 | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0102 | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0110 | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-10   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-11   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-20   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-21   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-22   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-23   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-24   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-25   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-26   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-27   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-28   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-30   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-31   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-32   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-33   | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka-tools                             | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka_2.12                              | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | kafka_2.13                              | 3.4.1.7.1.9.9-1 |
| Kafka   | org.apache.kafka | trogdor                                 | 3.4.1.7.1.9.9-1 |

| Project | groupId         | artifactId                                        | version         |
|---------|-----------------|---------------------------------------------------|-----------------|
| Knox    | org.apache.knox | gateway-adapter                                   | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-admin-ui                                  | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-applications                              | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-cloud-bindings                            | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-demo-ldap                                 | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-demo-ldap-launcher                        | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-discovery-ambari                          | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-discovery-cm                              | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-docker                                    | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-i18n                                      | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-i18n-logging-log4j                        | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-i18n-logging-sl4j                         | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-performance-test                          | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-ha                               | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-common        | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-concat        | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-hadoop-groups | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-pseudo        | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-regex         | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-identity-assertion-switchcase    | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-jersey                           | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-rewrite                          | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-rewrite-common                   | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-hostmap-static      | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-inbound-query-param | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-rewrite-func-service-registry    | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-rewrite-step-encrypt-uri         | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-rewrite-step-secure-query        | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-security-authc-anon              | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-security-authz-acls              | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-security-authz-composite         | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-security-clientcert              | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-security-hadoopauth              | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-security-jwt                     | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-security-pac4j                   | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-security-preauth                 | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-security-shiro                   | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-provider-security-webappsec               | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-release                                   | 1.3.0.7.1.9.9-1 |

| Project | groupId         | artifactId                       | version         |
|---------|-----------------|----------------------------------|-----------------|
| Knox    | org.apache.knox | gateway-server                   | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-server-launcher          | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-server-xforwarded-filter | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-admin            | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-as               | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-definitions      | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-hashicorp-vault  | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-hbase            | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-health           | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-hive             | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-idbroker         | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-impala           | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-jkg              | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-knoxsso          | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-knoxsout         | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-knoxtoken        | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-livy             | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-metadata         | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-nifi             | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-nifi-registry    | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-remoteconfig     | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-rm               | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-session          | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-storm            | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-test             | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-tgs              | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-vault            | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-service-webhdfs          | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-shell                    | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-shell-launcher           | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-shell-release            | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-shell-samples            | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-spi                      | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-test                     | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-test-idbroker            | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-test-release-utils       | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-test-utils               | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-topology-hadoop-xml      | 1.3.0.7.1.9.9-1 |
| Knox    | org.apache.knox | gateway-topology-simple          | 1.3.0.7.1.9.9-1 |

| Project | groupId           | artifactId                  | version          |
|---------|-------------------|-----------------------------|------------------|
| Knox    | org.apache.knox   | gateway-util-common         | 1.3.0.7.1.9.9-1  |
| Knox    | org.apache.knox   | gateway-util-configinjector | 1.3.0.7.1.9.9-1  |
| Knox    | org.apache.knox   | gateway-util-launcher       | 1.3.0.7.1.9.9-1  |
| Knox    | org.apache.knox   | gateway-util-urltemplate    | 1.3.0.7.1.9.9-1  |
| Knox    | org.apache.knox   | hadoop-examples             | 1.3.0.7.1.9.9-1  |
| Knox    | org.apache.knox   | knox-cli-launcher           | 1.3.0.7.1.9.9-1  |
| Knox    | org.apache.knox   | knox-homepage-ui            | 1.3.0.7.1.9.9-1  |
| Knox    | org.apache.knox   | knox-token-generation-ui    | 1.3.0.7.1.9.9-1  |
| Knox    | org.apache.knox   | knox-token-management-ui    | 1.3.0.7.1.9.9-1  |
| Knox    | org.apache.knox   | webhdfs-kerb-test           | 1.3.0.7.1.9.9-1  |
| Knox    | org.apache.knox   | webhdfs-test                | 1.3.0.7.1.9.9-1  |
| Kudu    | org.apache.kudu   | kudu-backup-tools           | 1.17.0.7.1.9.9-1 |
| Kudu    | org.apache.kudu   | kudu-backup2_2.11           | 1.17.0.7.1.9.9-1 |
| Kudu    | org.apache.kudu   | kudu-backup3_2.12           | 1.17.0.7.1.9.9-1 |
| Kudu    | org.apache.kudu   | kudu-client                 | 1.17.0.7.1.9.9-1 |
| Kudu    | org.apache.kudu   | kudu-hive                   | 1.17.0.7.1.9.9-1 |
| Kudu    | org.apache.kudu   | kudu-spark2-tools_2.11      | 1.17.0.7.1.9.9-1 |
| Kudu    | org.apache.kudu   | kudu-spark2_2.11            | 1.17.0.7.1.9.9-1 |
| Kudu    | org.apache.kudu   | kudu-spark3-tools_2.12      | 1.17.0.7.1.9.9-1 |
| Kudu    | org.apache.kudu   | kudu-spark3_2.12            | 1.17.0.7.1.9.9-1 |
| Kudu    | org.apache.kudu   | kudu-test-utils             | 1.17.0.7.1.9.9-1 |
| Livy    | org.apache.livy   | livy-api                    | 0.7.2.7.1.9.9-1  |
| Livy    | org.apache.livy   | livy-client-common          | 0.7.2.7.1.9.9-1  |
| Livy    | org.apache.livy   | livy-client-http            | 0.7.2.7.1.9.9-1  |
| Livy    | org.apache.livy   | livy-core_2.11              | 0.7.2.7.1.9.9-1  |
| Livy    | org.apache.livy   | livy-examples               | 0.7.2.7.1.9.9-1  |
| Livy    | org.apache.livy   | livy-integration-test       | 0.7.2.7.1.9.9-1  |
| Livy    | org.apache.livy   | livy-repl_2.11              | 0.7.2.7.1.9.9-1  |
| Livy    | org.apache.livy   | livy-rsc                    | 0.7.2.7.1.9.9-1  |
| Livy    | org.apache.livy   | livy-scala-api_2.11         | 0.7.2.7.1.9.9-1  |
| Livy    | org.apache.livy   | livy-server                 | 0.7.2.7.1.9.9-1  |
| Livy    | org.apache.livy   | livy-test-lib               | 0.7.2.7.1.9.9-1  |
| Livy    | org.apache.livy   | livy-thriftserver           | 0.7.2.7.1.9.9-1  |
| Livy    | org.apache.livy   | livy-thriftserver-session   | 0.7.2.7.1.9.9-1  |
| Lucene  | org.apache.lucene | lucene-analyzers-common     | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-analyzers-icu        | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-analyzers-kuromoji   | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-analyzers-morfologik | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-analyzers-nori       | 8.11.2.7.1.9.9-1 |

| Project | groupId           | artifactId                | version          |
|---------|-------------------|---------------------------|------------------|
| Lucene  | org.apache.lucene | lucene-analyzers-openssl  | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-analyzers-phonetic | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-analyzers-smartcn  | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-analyzers-stempel  | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-backward-codecs    | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-benchmark          | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-classification     | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-codecs             | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-core               | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-demo               | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-expressions        | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-facet              | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-grouping           | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-highlighter        | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-join               | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-memory             | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-misc               | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-monitor            | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-queries            | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-queryparser        | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-replicator         | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-sandbox            | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-spatial-extras     | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-spatial3d          | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-suggest            | 8.11.2.7.1.9.9-1 |
| Lucene  | org.apache.lucene | lucene-test-framework     | 8.11.2.7.1.9.9-1 |
| Oozie   | org.apache.oozie  | oozie-client              | 5.1.0.7.1.9.9-1  |
| Oozie   | org.apache.oozie  | oozie-core                | 5.1.0.7.1.9.9-1  |
| Oozie   | org.apache.oozie  | oozie-distro              | 5.1.0.7.1.9.9-1  |
| Oozie   | org.apache.oozie  | oozie-examples            | 5.1.0.7.1.9.9-1  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-api      | 5.1.0.7.1.9.9-1  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-client   | 5.1.0.7.1.9.9-1  |
| Oozie   | org.apache.oozie  | oozie-server              | 5.1.0.7.1.9.9-1  |
| Oozie   | org.apache.oozie  | oozie-sharelib-distcp     | 5.1.0.7.1.9.9-1  |
| Oozie   | org.apache.oozie  | oozie-sharelib-git        | 5.1.0.7.1.9.9-1  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hcatalog   | 5.1.0.7.1.9.9-1  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hive       | 5.1.0.7.1.9.9-1  |
| Oozie   | org.apache.oozie  | oozie-sharelib-hive2      | 5.1.0.7.1.9.9-1  |
| Oozie   | org.apache.oozie  | oozie-sharelib-oozie      | 5.1.0.7.1.9.9-1  |

| Project | groupId          | artifactId                     | version         |
|---------|------------------|--------------------------------|-----------------|
| Oozie   | org.apache.oozie | oozie-sharelib-spark           | 5.1.0.7.1.9.9-1 |
| Oozie   | org.apache.oozie | oozie-sharelib-spark3          | 5.1.0.7.1.9.9-1 |
| Oozie   | org.apache.oozie | oozie-sharelib-sqoop           | 5.1.0.7.1.9.9-1 |
| Oozie   | org.apache.oozie | oozie-sharelib-streaming       | 5.1.0.7.1.9.9-1 |
| Oozie   | org.apache.oozie | oozie-tools                    | 5.1.0.7.1.9.9-1 |
| Oozie   | org.apache.oozie | oozie-zookeeper-security-tests | 5.1.0.7.1.9.9-1 |
| ORC     | org.apache.orc   | orc-core                       | 1.5.1.7.1.9.9-1 |
| ORC     | org.apache.orc   | orc-examples                   | 1.5.1.7.1.9.9-1 |
| ORC     | org.apache.orc   | orc-mapreduce                  | 1.5.1.7.1.9.9-1 |
| ORC     | org.apache.orc   | orc-shims                      | 1.5.1.7.1.9.9-1 |
| ORC     | org.apache.orc   | orc-tools                      | 1.5.1.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-annotation-processing     | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-client                    | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-common                    | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-config                    | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-container-service         | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-docs                      | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-erasurecode               | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-client  | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-server  | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-test    | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-interface-admin           | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-interface-client          | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-interface-server          | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-managed-rocksdb           | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-rocks-native              | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-server-framework          | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-server-scm                | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-test-utils                | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | hdds-tools                     | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | mini-chaos-tests               | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | ozone-client                   | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | ozone-common                   | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | ozone-csi                      | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | ozone-datanode                 | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | ozone-dist                     | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | ozone-filesystem               | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | ozone-filesystem-common        | 1.4.0.7.1.9.9-1 |
| Ozone   | org.apache.ozone | ozone-filesystem-hadoop2       | 1.4.0.7.1.9.9-1 |

| Project | groupId            | artifactId                         | version           |
|---------|--------------------|------------------------------------|-------------------|
| Ozone   | org.apache.ozone   | ozone-filessystem-hadoop3          | 1.4.0.7.1.9.9-1   |
| Ozone   | org.apache.ozone   | ozone-filessystem-shaded           | 1.4.0.7.1.9.9-1   |
| Ozone   | org.apache.ozone   | ozone-httpgateway                  | 1.4.0.7.1.9.9-1   |
| Ozone   | org.apache.ozone   | ozone-insight                      | 1.4.0.7.1.9.9-1   |
| Ozone   | org.apache.ozone   | ozone-integration-test             | 1.4.0.7.1.9.9-1   |
| Ozone   | org.apache.ozone   | ozone-interface-client             | 1.4.0.7.1.9.9-1   |
| Ozone   | org.apache.ozone   | ozone-interface-storage            | 1.4.0.7.1.9.9-1   |
| Ozone   | org.apache.ozone   | ozone-manager                      | 1.4.0.7.1.9.9-1   |
| Ozone   | org.apache.ozone   | ozone-network-tests                | 1.4.0.7.1.9.9-1   |
| Ozone   | org.apache.ozone   | ozone-recon                        | 1.4.0.7.1.9.9-1   |
| Ozone   | org.apache.ozone   | ozone-reconcodegen                 | 1.4.0.7.1.9.9-1   |
| Ozone   | org.apache.ozone   | ozone-s3-secret-store              | 1.4.0.7.1.9.9-1   |
| Ozone   | org.apache.ozone   | ozone-s3gateway                    | 1.4.0.7.1.9.9-1   |
| Ozone   | org.apache.ozone   | ozone-tools                        | 1.4.0.7.1.9.9-1   |
| Ozone   | org.apache.ozone   | rocksdb-checkpoint-differ          | 1.4.0.7.1.9.9-1   |
| Parquet | org.apache.parquet | parquet-avro                       | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-cascading                  | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-cascading3                 | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-column                     | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-common                     | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-encoding                   | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-format-structures          | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-generator                  | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-hadoop                     | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-hadoop-bundle              | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-jackson                    | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-pig                        | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-pig-bundle                 | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-protobuf                   | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-scala_2.10                 | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-thrift                     | 1.10.99.7.1.9.9-1 |
| Parquet | org.apache.parquet | parquet-tools                      | 1.10.99.7.1.9.9-1 |
| Phoenix | org.apache.phoenix | phoenix-client-embedded-hbase-2.4  | 5.1.1.7.1.9.9-1   |
| Phoenix | org.apache.phoenix | phoenix-client-hbase-2.4           | 5.1.1.7.1.9.9-1   |
| Phoenix | org.apache.phoenix | phoenix-connectors-phoenix5-compat | 6.0.0.7.1.9.9-1   |
| Phoenix | org.apache.phoenix | phoenix-core                       | 5.1.1.7.1.9.9-1   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compat-2.1.6         | 5.1.1.7.1.9.9-1   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compat-2.2.5         | 5.1.1.7.1.9.9-1   |
| Phoenix | org.apache.phoenix | phoenix-hbase-compat-2.3.0         | 5.1.1.7.1.9.9-1   |



| Project | groupId            | artifactId                        | version         |
|---------|--------------------|-----------------------------------|-----------------|
| Phoenix | org.apache.phoenix | phoenix-hbase-compat-2.4.0        | 5.1.1.7.1.9.9-1 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compat-2.4.1        | 5.1.1.7.1.9.9-1 |
| Phoenix | org.apache.phoenix | phoenix-pherf                     | 5.1.1.7.1.9.9-1 |
| Phoenix | org.apache.phoenix | phoenix-queryserver               | 6.0.0.7.1.9.9-1 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-client        | 6.0.0.7.1.9.9-1 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-it            | 6.0.0.7.1.9.9-1 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-load-balancer | 6.0.0.7.1.9.9-1 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-orchestrator  | 6.0.0.7.1.9.9-1 |
| Phoenix | org.apache.phoenix | phoenix-server-hbase-2.4          | 5.1.1.7.1.9.9-1 |
| Phoenix | org.apache.phoenix | phoenix-tracing-webapp            | 5.1.1.7.1.9.9-1 |
| Phoenix | org.apache.phoenix | phoenix5-hive                     | 6.0.0.7.1.9.9-1 |
| Phoenix | org.apache.phoenix | phoenix5-hive-shaded              | 6.0.0.7.1.9.9-1 |
| Phoenix | org.apache.phoenix | phoenix5-spark                    | 6.0.0.7.1.9.9-1 |
| Phoenix | org.apache.phoenix | phoenix5-spark-shaded             | 6.0.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | conditions-enrichers              | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | credentialbuilder                 | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | embeddedwebserver                 | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | jisql                             | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ldapconfigcheck                   | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-adls-plugin                | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-atlas-plugin               | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-atlas-plugin-shim          | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-authn                      | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-common-ha                  | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-distro                     | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-examples-distro            | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-hbase-plugin               | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-hbase-plugin-shim          | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-hdfs-plugin                | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-hdfs-plugin-shim           | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-hive-plugin                | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-hive-plugin-shim           | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-intg                       | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-kafka-connect-plugin       | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-kafka-plugin               | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-kafka-plugin-shim          | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-kms                        | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-kms-plugin                 | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger  | ranger-kms-plugin-shim            | 2.4.0.7.1.9.9-1 |

| Project | groupId           | artifactId                    | version         |
|---------|-------------------|-------------------------------|-----------------|
| Ranger  | org.apache.ranger | ranger-knox-plugin            | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-knox-plugin-shim       | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-kudu-plugin            | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-kylin-plugin           | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-kylin-plugin-shim      | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-metrics                | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-nifi-plugin            | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-nifi-registry-plugin   | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-ozone-plugin           | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-ozone-plugin-shim      | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-plugin-classloader     | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-plugins-audit          | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-plugins-common         | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-plugins-cred           | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-plugins-installer      | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-policymigration        | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-raz-adls               | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-raz-chained-plugins    | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-raz-hook-abfs          | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-raz-intg               | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-raz-processor          | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-rms-common             | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-rms-hive               | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-rms-plugins-common     | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-rms-tools              | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-rms-webapp             | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-sampleapp-plugin       | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-schema-registry-plugin | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-solr-plugin            | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-solr-plugin-shim       | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin           | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin-shim      | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-storm-plugin           | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-storm-plugin-shim      | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-tagsync                | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-tools                  | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-util                   | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-yarn-plugin            | 2.4.0.7.1.9.9-1 |
| Ranger  | org.apache.ranger | ranger-yarn-plugin-shim       | 2.4.0.7.1.9.9-1 |

| Project | groupId           | artifactId                     | version          |
|---------|-------------------|--------------------------------|------------------|
| Ranger  | org.apache.ranger | sample-client                  | 2.4.0.7.1.9.9-1  |
| Ranger  | org.apache.ranger | sampleapp                      | 2.4.0.7.1.9.9-1  |
| Ranger  | org.apache.ranger | shaded-raz-hook-abfs           | 2.4.0.7.1.9.9-1  |
| Ranger  | org.apache.ranger | ugsync-util                    | 2.4.0.7.1.9.9-1  |
| Ranger  | org.apache.ranger | unixauthclient                 | 2.4.0.7.1.9.9-1  |
| Ranger  | org.apache.ranger | unixauthservice                | 2.4.0.7.1.9.9-1  |
| Ranger  | org.apache.ranger | unixusersync                   | 2.4.0.7.1.9.9-1  |
| Solr    | org.apache.solr   | solr-analysis-extras           | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-analytics                 | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-cell                      | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-core                      | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-dataimporthandler         | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-dataimporthandler-extras  | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-gcs-repository            | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-jaegertracer-configurator | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-langid                    | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-ltr                       | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-prometheus-exporter       | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-s3-repository             | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-security-util             | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-solrj                     | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-test-framework            | 8.11.2.7.1.9.9-1 |
| Solr    | org.apache.solr   | solr-velocity                  | 8.11.2.7.1.9.9-1 |
| Spark   | org.apache.spark  | spark-avro_2.11                | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-catalyst_2.11            | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-core_2.11                | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-graphx_2.11              | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-hadoop-cloud_2.11        | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-hive_2.11                | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-kubernetes_2.11          | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-kvstore_2.11             | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-launcher_2.11            | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-mllib-local_2.11         | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-mllib_2.11               | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-network-common_2.11      | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-network-shuffle_2.11     | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-network-yarn_2.11        | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-repl_2.11                | 2.4.8.7.1.9.9-1  |
| Spark   | org.apache.spark  | spark-sketch_2.11              | 2.4.8.7.1.9.9-1  |

| Project  | groupId             | artifactId                               | version         |
|----------|---------------------|------------------------------------------|-----------------|
| Spark    | org.apache.spark    | spark-sql-kafka-0-10_2.11                | 2.4.8.7.1.9.9-1 |
| Spark    | org.apache.spark    | spark-sql_2.11                           | 2.4.8.7.1.9.9-1 |
| Spark    | org.apache.spark    | spark-streaming-kafka-0-10-assembly_2.11 | 2.4.8.7.1.9.9-1 |
| Spark    | org.apache.spark    | spark-streaming-kafka-0-10_2.11          | 2.4.8.7.1.9.9-1 |
| Spark    | org.apache.spark    | spark-streaming_2.11                     | 2.4.8.7.1.9.9-1 |
| Spark    | org.apache.spark    | spark-tags_2.11                          | 2.4.8.7.1.9.9-1 |
| Spark    | org.apache.spark    | spark-token-provider-kafka-0-10_2.11     | 2.4.8.7.1.9.9-1 |
| Spark    | org.apache.spark    | spark-unsafe_2.11                        | 2.4.8.7.1.9.9-1 |
| Spark    | org.apache.spark    | spark-yarn_2.11                          | 2.4.8.7.1.9.9-1 |
| Sqoop    | org.apache.sqoop    | sqoop                                    | 1.4.7.7.1.9.9-1 |
| Sqoop    | org.apache.sqoop    | sqoop-test                               | 1.4.7.7.1.9.9-1 |
| Tez      | org.apache.tez      | hadoop-shim                              | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | hadoop-shim-2.8                          | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-api                                  | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-aux-services                         | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-common                               | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-dag                                  | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-examples                             | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-ext-service-tests                    | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-history-parser                       | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-javadoc-tools                        | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-job-analyzer                         | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-mapreduce                            | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-protobuf-history-plugin              | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-runtime-internals                    | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-runtime-library                      | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-tests                                | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-yarn-timeline-cache-plugin           | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-yarn-timeline-history                | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-yarn-timeline-history-with-acls      | 0.9.1.7.1.9.9-1 |
| Tez      | org.apache.tez      | tez-yarn-timeline-history-with-fs        | 0.9.1.7.1.9.9-1 |
| Zeppelin | org.apache.zeppelin | zeppelin-angular                         | 0.8.2.7.1.9.9-1 |
| Zeppelin | org.apache.zeppelin | zeppelin-display                         | 0.8.2.7.1.9.9-1 |
| Zeppelin | org.apache.zeppelin | zeppelin-interpreter                     | 0.8.2.7.1.9.9-1 |
| Zeppelin | org.apache.zeppelin | zeppelin-jdbc                            | 0.8.2.7.1.9.9-1 |
| Zeppelin | org.apache.zeppelin | zeppelin-jupyter                         | 0.8.2.7.1.9.9-1 |
| Zeppelin | org.apache.zeppelin | zeppelin-livy                            | 0.8.2.7.1.9.9-1 |
| Zeppelin | org.apache.zeppelin | zeppelin-markdown                        | 0.8.2.7.1.9.9-1 |
| Zeppelin | org.apache.zeppelin | zeppelin-server                          | 0.8.2.7.1.9.9-1 |

| Project   | groupId              | artifactId                   | version         |
|-----------|----------------------|------------------------------|-----------------|
| Zeppelin  | org.apache.zeppelin  | zeppelin-shell               | 0.8.2.7.1.9.9-1 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-zengine             | 0.8.2.7.1.9.9-1 |
| ZooKeeper | org.apache.zookeeper | zookeeper                    | 3.8.1.7.1.9.9-1 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-fatjar     | 3.8.1.7.1.9.9-1 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-loggraph   | 3.8.1.7.1.9.9-1 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-rest       | 3.8.1.7.1.9.9-1 |
| ZooKeeper | org.apache.zookeeper | zookeeper-it                 | 3.8.1.7.1.9.9-1 |
| ZooKeeper | org.apache.zookeeper | zookeeper-jute               | 3.8.1.7.1.9.9-1 |
| ZooKeeper | org.apache.zookeeper | zookeeper-prometheus-metrics | 3.8.1.7.1.9.9-1 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-election   | 3.8.1.7.1.9.9-1 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-queue      | 3.8.1.7.1.9.9-1 |

## Runtime 7.1.9.14-2

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          | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-aws-s3-bridge          | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-azure-adls-bridge      | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-classification-updater | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-client-common          | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-client-v1              | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-client-v2              | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-common                 | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-distro                 | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-docs                   | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-graphdb-api            | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-graphdb-common         | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-graphdb-janus          | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-hdfs-bridge            | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-index-repair-tool      | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-intg                   | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-janusgraph-hbase2      | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-notification           | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-plugin-classloader     | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-repository             | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-server-api             | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | atlas-testtools              | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | hbase-bridge                 | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | hbase-bridge-shim            | 3.0.0.7.1.9.14-2 |
| Atlas   | org.apache.atlas | hbase-testing-util           | 3.0.0.7.1.9.14-2 |

| Project | groupId            | artifactId                   | version           |
|---------|--------------------|------------------------------|-------------------|
| Atlas   | org.apache.atlas   | hdfs-model                   | 3.0.0.7.1.9.14-2  |
| Atlas   | org.apache.atlas   | hive-bridge                  | 3.0.0.7.1.9.14-2  |
| Atlas   | org.apache.atlas   | hive-bridge-shim             | 3.0.0.7.1.9.14-2  |
| Atlas   | org.apache.atlas   | impala-bridge                | 3.0.0.7.1.9.14-2  |
| Atlas   | org.apache.atlas   | impala-bridge-shim           | 3.0.0.7.1.9.14-2  |
| Atlas   | org.apache.atlas   | impala-hook-api              | 3.0.0.7.1.9.14-2  |
| Atlas   | org.apache.atlas   | kafka-bridge                 | 3.0.0.7.1.9.14-2  |
| Atlas   | org.apache.atlas   | kafka-bridge-shim            | 3.0.0.7.1.9.14-2  |
| Atlas   | org.apache.atlas   | navigator-to-atlas           | 3.0.0.7.1.9.14-2  |
| Atlas   | org.apache.atlas   | sample-app                   | 3.0.0.7.1.9.14-2  |
| Atlas   | org.apache.atlas   | sqoop-bridge                 | 3.0.0.7.1.9.14-2  |
| Atlas   | org.apache.atlas   | sqoop-bridge-shim            | 3.0.0.7.1.9.14-2  |
| Avro    | org.apache.avro    | avro                         | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-android                 | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-codegen-test            | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-compiler                | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-grpc                    | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-ipc                     | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-ipc-jetty               | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-ipc-netty               | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-mapred                  | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-maven-plugin            | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-perf                    | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-protobuf                | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-service-archetype       | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-test-custom-conversions | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-thrift                  | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | avro-tools                   | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | trevni-avro                  | 1.11.1.7.1.9.14-2 |
| Avro    | org.apache.avro    | trevni-core                  | 1.11.1.7.1.9.14-2 |
| Calcite | org.apache.calcite | calcite-babel                | 1.19.0.7.1.9.14-2 |
| Calcite | org.apache.calcite | calcite-core                 | 1.19.0.7.1.9.14-2 |
| Calcite | org.apache.calcite | calcite-druid                | 1.19.0.7.1.9.14-2 |
| Calcite | org.apache.calcite | calcite-linq4j               | 1.19.0.7.1.9.14-2 |
| Calcite | org.apache.calcite | calcite-server               | 1.19.0.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop  | hadoop-aliyun                | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop  | hadoop-annotations           | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop  | hadoop-archive-logs          | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop  | hadoop-archives              | 3.1.1.7.1.9.14-2  |

| Project | groupId           | artifactId                         | version          |
|---------|-------------------|------------------------------------|------------------|
| Hadoop  | org.apache.hadoop | hadoop-assemblies                  | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-auth                        | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-aws                         | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-azure                       | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-azure-datalake              | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-build-tools                 | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-client                      | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-client-api                  | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-client-integration-tests    | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-client-miniclust            | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-client-runtime              | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-cloud-storage               | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-common                      | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-datajoin                    | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-distcp                      | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-extras                      | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-fs2img                      | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-gridmix                     | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs                        | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-client                 | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-httpfs                 | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-native-client          | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-nfs                    | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-hdfs-rbf                    | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-kafka                       | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-kms                         | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-app        | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-common     | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-core       | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-hs         | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-hs-plugins | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-jobclient  | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-nativetask | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-shuffle    | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-client-uploader   | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-mapreduce-examples          | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-maven-plugins               | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-miniclust                   | 3.1.1.7.1.9.14-2 |
| Hadoop  | org.apache.hadoop | hadoop-minikdc                     | 3.1.1.7.1.9.14-2 |

| Project | groupId           | artifactId                                        | version           |
|---------|-------------------|---------------------------------------------------|-------------------|
| Hadoop  | org.apache.hadoop | hadoop-nfs                                        | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-openstack                                  | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-resourceestimator                          | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-rumen                                      | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-sls                                        | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-streaming                                  | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-tools-dist                                 | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-api                                   | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-distributedshell         | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-applications-unmanaged-am-launcher    | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-client                                | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-common                                | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-registry                              | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-applicationhistoryservice      | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-common                         | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-nodemanager                    | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-resourcemanager                | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-router                         | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-sharedcachemanager             | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-tests                          | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timeline-pluginstorage         | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice                | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-client   | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-common   | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-server-2 | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-timelineservice-hbase-tests    | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-server-web-proxy                      | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-api                          | 3.1.1.7.1.9.14-2  |
| Hadoop  | org.apache.hadoop | hadoop-yarn-services-core                         | 3.1.1.7.1.9.14-2  |
| HBase   | org.apache.hbase  | hbase-annotations                                 | 2.4.17.7.1.9.14-2 |
| HBase   | org.apache.hbase  | hbase-asyncfs                                     | 2.4.17.7.1.9.14-2 |
| HBase   | org.apache.hbase  | hbase-checkstyle                                  | 2.4.17.7.1.9.14-2 |
| HBase   | org.apache.hbase  | hbase-client                                      | 2.4.17.7.1.9.14-2 |
| HBase   | org.apache.hbase  | hbase-client-project                              | 2.4.17.7.1.9.14-2 |
| HBase   | org.apache.hbase  | hbase-common                                      | 2.4.17.7.1.9.14-2 |
| HBase   | org.apache.hbase  | hbase-endpoint                                    | 2.4.17.7.1.9.14-2 |
| HBase   | org.apache.hbase  | hbase-examples                                    | 2.4.17.7.1.9.14-2 |
| HBase   | org.apache.hbase  | hbase-external-blockcache                         | 2.4.17.7.1.9.14-2 |
| HBase   | org.apache.hbase  | hbase-hadoop-compat                               | 2.4.17.7.1.9.14-2 |



| Project | groupId          | artifactId                       | version             |
|---------|------------------|----------------------------------|---------------------|
| HBase   | org.apache.hbase | hbase-hadoop2-compat             | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-hbtop                      | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-http                       | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-it                         | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-logging                    | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-mapreduce                  | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-metrics                    | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-metrics-api                | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-procedure                  | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-protocol                   | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-protocol-shaded            | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-replication                | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-resource-bundle            | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-rest                       | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-rsgroup                    | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-server                     | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-shaded-client              | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-shaded-client-byo-hadoop   | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-shaded-client-project      | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-shaded-mapreduce           | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-shaded-testing-util        | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-shaded-testing-util-tester | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-shell                      | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-testing-util               | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-thrift                     | 2.4.17.7.1.9.14-2   |
| HBase   | org.apache.hbase | hbase-zookeeper                  | 2.4.17.7.1.9.14-2   |
| Hive    | org.apache.hive  | hive-beeline                     | 3.1.3000.7.1.9.14-2 |
| Hive    | org.apache.hive  | hive-blobstore                   | 3.1.3000.7.1.9.14-2 |
| Hive    | org.apache.hive  | hive-classification              | 3.1.3000.7.1.9.14-2 |
| Hive    | org.apache.hive  | hive-cli                         | 3.1.3000.7.1.9.14-2 |
| Hive    | org.apache.hive  | hive-common                      | 3.1.3000.7.1.9.14-2 |
| Hive    | org.apache.hive  | hive-contrib                     | 3.1.3000.7.1.9.14-2 |
| Hive    | org.apache.hive  | hive-exec                        | 3.1.3000.7.1.9.14-2 |
| Hive    | org.apache.hive  | hive-hbase-handler               | 3.1.3000.7.1.9.14-2 |
| Hive    | org.apache.hive  | hive-hcatalog-it-unit            | 3.1.3000.7.1.9.14-2 |
| Hive    | org.apache.hive  | hive-hplsql                      | 3.1.3000.7.1.9.14-2 |
| Hive    | org.apache.hive  | hive-it-custom-serde             | 3.1.3000.7.1.9.14-2 |
| Hive    | org.apache.hive  | hive-it-minikdc                  | 3.1.3000.7.1.9.14-2 |
| Hive    | org.apache.hive  | hive-it-qfile                    | 3.1.3000.7.1.9.14-2 |

| Project                  | groupId              | artifactId                               | version             |
|--------------------------|----------------------|------------------------------------------|---------------------|
| Hive                     | org.apache.hive      | hive-it-qfile-kudu                       | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-it-test-serde                       | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-it-unit                             | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-it-unit-hadoop2                     | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-it-util                             | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-jdbc                                | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-jdbc-handler                        | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-jmh                                 | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-kryo-registrator                    | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-kudu-handler                        | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-llap-client                         | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-llap-common                         | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-llap-ext-client                     | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-llap-server                         | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-llap-tez                            | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-metastore                           | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-pre-upgrade                         | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-serde                               | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-service                             | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-service-rpc                         | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-shims                               | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-spark-client                        | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-standalone-metastore                | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-storage-api                         | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-streaming                           | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-testutils                           | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | hive-vector-code-gen                     | 3.1.3000.7.1.9.14-2 |
| Hive                     | org.apache.hive      | kafka-handler                            | 3.1.3000.7.1.9.14-2 |
| Hive Warehouse Connector | com.hortonworks.hive | hive-warehouse-connector_2.11            | 1.0.0.7.1.9.14-2    |
| Kafka                    | org.apache.kafka     | ci                                       | 3.4.1.7.1.9.14-2    |
| Kafka                    | org.apache.kafka     | connect                                  | 3.4.1.7.1.9.14-2    |
| Kafka                    | org.apache.kafka     | connect-api                              | 3.4.1.7.1.9.14-2    |
| Kafka                    | org.apache.kafka     | connect-basic-auth-extension             | 3.4.1.7.1.9.14-2    |
| Kafka                    | org.apache.kafka     | connect-cloudera-authorization-extension | 3.4.1.7.1.9.14-2    |
| Kafka                    | org.apache.kafka     | connect-cloudera-common                  | 3.4.1.7.1.9.14-2    |
| Kafka                    | org.apache.kafka     | connect-cloudera-secret-storage          | 3.4.1.7.1.9.14-2    |
| Kafka                    | org.apache.kafka     | connect-cloudera-security-policies       | 3.4.1.7.1.9.14-2    |
| Kafka                    | org.apache.kafka     | connect-file                             | 3.4.1.7.1.9.14-2    |
| Kafka                    | org.apache.kafka     | connect-json                             | 3.4.1.7.1.9.14-2    |

| Project | groupId          | artifactId                              | version          |
|---------|------------------|-----------------------------------------|------------------|
| Kafka   | org.apache.kafka | connect-mirror                          | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | connect-mirror-client                   | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | connect-runtime                         | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | connect-transforms                      | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | generator                               | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-clients                           | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-cloudera-metrics-reporter_2.12    | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-cloudera-metrics-reporter_2.13    | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-cloudera-plugins                  | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-examples                          | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-group-coordinator                 | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-log4j-appender                    | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-metadata                          | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-raft                              | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-server-common                     | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-shell                             | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-storage                           | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-storage-api                       | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams                           | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-examples                  | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-scala_2.12                | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-scala_2.13                | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-test-utils                | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0100 | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0101 | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0102 | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-0110 | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-10   | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-11   | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-20   | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-21   | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-22   | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-23   | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-24   | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-25   | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-26   | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-27   | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-28   | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-30   | 3.4.1.7.1.9.14-2 |

| Project | groupId          | artifactId                                        | version          |
|---------|------------------|---------------------------------------------------|------------------|
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-31             | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-32             | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-streams-upgrade-system-tests-33             | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka-tools                                       | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka_2.12                                        | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | kafka_2.13                                        | 3.4.1.7.1.9.14-2 |
| Kafka   | org.apache.kafka | trogdor                                           | 3.4.1.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-adapter                                   | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-admin-ui                                  | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-applications                              | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-cloud-bindings                            | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-demo-ldap                                 | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-demo-ldap-launcher                        | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-discovery-ambari                          | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-discovery-cm                              | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-docker                                    | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-i18n                                      | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-i18n-logging-log4j                        | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-i18n-logging-sl4j                         | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-performance-test                          | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-ha                               | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-common        | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-concat        | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-hadoop-groups | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-pseudo        | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-regex         | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-identity-assertion-switchcase    | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-jersey                           | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-rewrite                          | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-common                   | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-func-hostmap-static      | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-func-inbound-query-param | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-func-service-registry    | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-step-encrypt-uri         | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-rewrite-step-secure-query        | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-security-authc-anon              | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-security-authz-acls              | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-security-authz-composite         | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox  | gateway-provider-security-clientcert              | 1.3.0.7.1.9.14-2 |

| Project | groupId         | artifactId                           | version          |
|---------|-----------------|--------------------------------------|------------------|
| Knox    | org.apache.knox | gateway-provider-security-hadoopauth | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-provider-security-jwt        | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-provider-security-pac4j      | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-provider-security-preauth    | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-provider-security-shiro      | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-provider-security-webappsec  | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-release                      | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-server                       | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-server-launcher              | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-server-xforwarded-filter     | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-admin                | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-as                   | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-definitions          | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-hashicorp-vault      | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-hbase                | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-health               | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-hive                 | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-idbroker             | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-impala               | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-jkg                  | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-knoxssso             | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-knoxssout            | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-knoxtoken            | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-livy                 | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-metadata             | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-nifi                 | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-nifi-registry        | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-remoteconfig         | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-rm                   | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-session              | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-storm                | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-test                 | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-tgs                  | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-vault                | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-service-webhdfs              | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-shell                        | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-shell-launcher               | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-shell-release                | 1.3.0.7.1.9.14-2 |
| Knox    | org.apache.knox | gateway-shell-samples                | 1.3.0.7.1.9.14-2 |

| Project | groupId         | artifactId                  | version           |
|---------|-----------------|-----------------------------|-------------------|
| Knox    | org.apache.knox | gateway-spi                 | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | gateway-test                | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | gateway-test-idbroker       | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | gateway-test-release-utils  | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | gateway-test-utils          | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | gateway-topology-hadoop-xml | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | gateway-topology-simple     | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | gateway-util-common         | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | gateway-util-configinjector | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | gateway-util-launcher       | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | gateway-util-urltemplate    | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | hadoop-examples             | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | knox-cli-launcher           | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | knox-homepage-ui            | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | knox-token-generation-ui    | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | knox-token-management-ui    | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | webhdfs-kerb-test           | 1.3.0.7.1.9.14-2  |
| Knox    | org.apache.knox | webhdfs-test                | 1.3.0.7.1.9.14-2  |
| Kudu    | org.apache.kudu | kudu-backup-tools           | 1.17.0.7.1.9.14-2 |
| Kudu    | org.apache.kudu | kudu-backup2_2.11           | 1.17.0.7.1.9.14-2 |
| Kudu    | org.apache.kudu | kudu-backup3_2.12           | 1.17.0.7.1.9.14-2 |
| Kudu    | org.apache.kudu | kudu-client                 | 1.17.0.7.1.9.14-2 |
| Kudu    | org.apache.kudu | kudu-hive                   | 1.17.0.7.1.9.14-2 |
| Kudu    | org.apache.kudu | kudu-spark2-tools_2.11      | 1.17.0.7.1.9.14-2 |
| Kudu    | org.apache.kudu | kudu-spark2_2.11            | 1.17.0.7.1.9.14-2 |
| Kudu    | org.apache.kudu | kudu-spark3-tools_2.12      | 1.17.0.7.1.9.14-2 |
| Kudu    | org.apache.kudu | kudu-spark3_2.12            | 1.17.0.7.1.9.14-2 |
| Kudu    | org.apache.kudu | kudu-test-utils             | 1.17.0.7.1.9.14-2 |
| Livy    | org.apache.livy | livy-api                    | 0.7.2.7.1.9.14-2  |
| Livy    | org.apache.livy | livy-client-common          | 0.7.2.7.1.9.14-2  |
| Livy    | org.apache.livy | livy-client-http            | 0.7.2.7.1.9.14-2  |
| Livy    | org.apache.livy | livy-core_2.11              | 0.7.2.7.1.9.14-2  |
| Livy    | org.apache.livy | livy-examples               | 0.7.2.7.1.9.14-2  |
| Livy    | org.apache.livy | livy-integration-test       | 0.7.2.7.1.9.14-2  |
| Livy    | org.apache.livy | livy-repl_2.11              | 0.7.2.7.1.9.14-2  |
| Livy    | org.apache.livy | livy-rsc                    | 0.7.2.7.1.9.14-2  |
| Livy    | org.apache.livy | livy-scala-api_2.11         | 0.7.2.7.1.9.14-2  |
| Livy    | org.apache.livy | livy-server                 | 0.7.2.7.1.9.14-2  |
| Livy    | org.apache.livy | livy-test-lib               | 0.7.2.7.1.9.14-2  |

| Project | groupId           | artifactId                  | version           |
|---------|-------------------|-----------------------------|-------------------|
| Livy    | org.apache.livy   | livy-thriftserver           | 0.7.2.7.1.9.14-2  |
| Livy    | org.apache.livy   | livy-thriftserver-session   | 0.7.2.7.1.9.14-2  |
| Lucene  | org.apache.lucene | lucene-analyzers-common     | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-icu        | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-kuromoji   | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-morfologik | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-nori       | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-opennlp    | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-phonetic   | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-smartcn    | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-analyzers-stempel    | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-backward-codecs      | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-benchmark            | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-classification       | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-codecs               | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-core                 | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-demo                 | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-expressions          | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-facet                | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-grouping             | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-highlighter          | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-join                 | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-memory               | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-misc                 | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-monitor              | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-queries              | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-queryparser          | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-replicator           | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-sandbox              | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-spatial-extras       | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-spatial3d            | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-suggest              | 8.11.2.7.1.9.14-2 |
| Lucene  | org.apache.lucene | lucene-test-framework       | 8.11.2.7.1.9.14-2 |
| Oozie   | org.apache.oozie  | oozie-client                | 5.1.0.7.1.9.14-2  |
| Oozie   | org.apache.oozie  | oozie-core                  | 5.1.0.7.1.9.14-2  |
| Oozie   | org.apache.oozie  | oozie-distro                | 5.1.0.7.1.9.14-2  |
| Oozie   | org.apache.oozie  | oozie-examples              | 5.1.0.7.1.9.14-2  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-api        | 5.1.0.7.1.9.14-2  |
| Oozie   | org.apache.oozie  | oozie-fluent-job-client     | 5.1.0.7.1.9.14-2  |

| Project | groupId          | artifactId                     | version          |
|---------|------------------|--------------------------------|------------------|
| Oozie   | org.apache.oozie | oozie-server                   | 5.1.0.7.1.9.14-2 |
| Oozie   | org.apache.oozie | oozie-sharelib-distcp          | 5.1.0.7.1.9.14-2 |
| Oozie   | org.apache.oozie | oozie-sharelib-git             | 5.1.0.7.1.9.14-2 |
| Oozie   | org.apache.oozie | oozie-sharelib-hcatalog        | 5.1.0.7.1.9.14-2 |
| Oozie   | org.apache.oozie | oozie-sharelib-hive            | 5.1.0.7.1.9.14-2 |
| Oozie   | org.apache.oozie | oozie-sharelib-hive2           | 5.1.0.7.1.9.14-2 |
| Oozie   | org.apache.oozie | oozie-sharelib-oozie           | 5.1.0.7.1.9.14-2 |
| Oozie   | org.apache.oozie | oozie-sharelib-spark           | 5.1.0.7.1.9.14-2 |
| Oozie   | org.apache.oozie | oozie-sharelib-spark3          | 5.1.0.7.1.9.14-2 |
| Oozie   | org.apache.oozie | oozie-sharelib-sqoop           | 5.1.0.7.1.9.14-2 |
| Oozie   | org.apache.oozie | oozie-sharelib-streaming       | 5.1.0.7.1.9.14-2 |
| Oozie   | org.apache.oozie | oozie-tools                    | 5.1.0.7.1.9.14-2 |
| Oozie   | org.apache.oozie | oozie-zookeeper-security-tests | 5.1.0.7.1.9.14-2 |
| ORC     | org.apache.orc   | orc-core                       | 1.5.1.7.1.9.14-2 |
| ORC     | org.apache.orc   | orc-examples                   | 1.5.1.7.1.9.14-2 |
| ORC     | org.apache.orc   | orc-mapreduce                  | 1.5.1.7.1.9.14-2 |
| ORC     | org.apache.orc   | orc-shims                      | 1.5.1.7.1.9.14-2 |
| ORC     | org.apache.orc   | orc-tools                      | 1.5.1.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-annotation-processing     | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-client                    | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-common                    | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-config                    | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-container-service         | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-docs                      | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-erasurecode               | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-client  | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-server  | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-hadoop-dependency-test    | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-interface-admin           | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-interface-client          | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-interface-server          | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-managed-rocksdbs          | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-rocks-native              | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-server-framework          | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-server-scm                | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-test-utils                | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | hdds-tools                     | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | mini-chaos-tests               | 1.4.0.7.1.9.14-2 |
| Ozone   | org.apache.ozone | ozone-client                   | 1.4.0.7.1.9.14-2 |



| Project | groupId            | artifactId                | version            |
|---------|--------------------|---------------------------|--------------------|
| Ozone   | org.apache.ozone   | ozone-common              | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-csi                 | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-datanode            | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-dist                | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-filessystem         | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-filessystem-common  | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-filessystem-hadoop2 | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-filessystem-hadoop3 | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-filessystem-shaded  | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-httpfsgateway       | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-insight             | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-integration-test    | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-interface-client    | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-interface-storage   | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-manager             | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-network-tests       | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-recon               | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-reconcodegen        | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-s3-secret-store     | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-s3gateway           | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | ozone-tools               | 1.4.0.7.1.9.14-2   |
| Ozone   | org.apache.ozone   | rocksdb-checkpoint-differ | 1.4.0.7.1.9.14-2   |
| Parquet | org.apache.parquet | parquet-avro              | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-cascading         | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-cascading3        | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-column            | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-common            | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-encoding          | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-format-structures | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-generator         | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-hadoop            | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-hadoop-bundle     | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-jackson           | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-pig               | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-pig-bundle        | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-protobuf          | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-scala_2.10        | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-thrift            | 1.10.99.7.1.9.14-2 |
| Parquet | org.apache.parquet | parquet-tools             | 1.10.99.7.1.9.14-2 |

| Project | groupId            | artifactId                         | version          |
|---------|--------------------|------------------------------------|------------------|
| Phoenix | org.apache.phoenix | phoenix-client-embedded-hbase-2.4  | 5.1.1.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-client-hbase-2.4           | 5.1.1.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-connectors-phoenix5-compat | 6.0.0.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-core                       | 5.1.1.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compat-2.1.6         | 5.1.1.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compat-2.2.5         | 5.1.1.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compat-2.3.0         | 5.1.1.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compat-2.4.0         | 5.1.1.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-hbase-compat-2.4.1         | 5.1.1.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-pherf                      | 5.1.1.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-queryserver                | 6.0.0.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-client         | 6.0.0.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-it             | 6.0.0.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-load-balancer  | 6.0.0.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-queryserver-orchestrator   | 6.0.0.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-server-hbase-2.4           | 5.1.1.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix-tracing-webapp             | 5.1.1.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix5-hive                      | 6.0.0.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix5-hive-shaded               | 6.0.0.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix5-spark                     | 6.0.0.7.1.9.14-2 |
| Phoenix | org.apache.phoenix | phoenix5-spark-shaded              | 6.0.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | conditions-enrichers               | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | credentialbuilder                  | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | embeddedwebserver                  | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | jisql                              | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | ldapconfigcheck                    | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | ranger-adls-plugin                 | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | ranger-atlas-plugin                | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | ranger-atlas-plugin-shim           | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | ranger-authn                       | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | ranger-common-ha                   | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | ranger-distro                      | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | ranger-examples-distro             | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | ranger-hbase-plugin                | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | ranger-hbase-plugin-shim           | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | ranger-hdfs-plugin                 | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | ranger-hdfs-plugin-shim            | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | ranger-hive-plugin                 | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger  | ranger-hive-plugin-shim            | 2.4.0.7.1.9.14-2 |

| Project | groupId           | artifactId                    | version          |
|---------|-------------------|-------------------------------|------------------|
| Ranger  | org.apache.ranger | ranger-intg                   | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-kafka-connect-plugin   | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-kafka-plugin           | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-kafka-plugin-shim      | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-kms                    | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-kms-plugin             | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-kms-plugin-shim        | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-knox-plugin            | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-knox-plugin-shim       | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-kudu-plugin            | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-kylin-plugin           | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-kylin-plugin-shim      | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-metrics                | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-nifi-plugin            | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-nifi-registry-plugin   | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-ozone-plugin           | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-ozone-plugin-shim      | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-plugin-classloader     | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-plugins-audit          | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-plugins-common         | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-plugins-cred           | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-plugins-installer      | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-policymigration        | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-raz-adls               | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-raz-chained-plugins    | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-raz-hook-abfs          | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-raz-intg               | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-raz-processor          | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-rms-common             | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-rms-hive               | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-rms-plugins-common     | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-rms-tools              | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-rms-webapp             | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-sampleapp-plugin       | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-schema-registry-plugin | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-solr-plugin            | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-solr-plugin-shim       | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin           | 2.4.0.7.1.9.14-2 |
| Ranger  | org.apache.ranger | ranger-sqoop-plugin-shim      | 2.4.0.7.1.9.14-2 |

| Project | groupId           | artifactId                     | version           |
|---------|-------------------|--------------------------------|-------------------|
| Ranger  | org.apache.ranger | ranger-storm-plugin            | 2.4.0.7.1.9.14-2  |
| Ranger  | org.apache.ranger | ranger-storm-plugin-shim       | 2.4.0.7.1.9.14-2  |
| Ranger  | org.apache.ranger | ranger-tagsync                 | 2.4.0.7.1.9.14-2  |
| Ranger  | org.apache.ranger | ranger-tools                   | 2.4.0.7.1.9.14-2  |
| Ranger  | org.apache.ranger | ranger-util                    | 2.4.0.7.1.9.14-2  |
| Ranger  | org.apache.ranger | ranger-yarn-plugin             | 2.4.0.7.1.9.14-2  |
| Ranger  | org.apache.ranger | ranger-yarn-plugin-shim        | 2.4.0.7.1.9.14-2  |
| Ranger  | org.apache.ranger | sample-client                  | 2.4.0.7.1.9.14-2  |
| Ranger  | org.apache.ranger | sampleapp                      | 2.4.0.7.1.9.14-2  |
| Ranger  | org.apache.ranger | shaded-raz-hook-abfs           | 2.4.0.7.1.9.14-2  |
| Ranger  | org.apache.ranger | ugsync-util                    | 2.4.0.7.1.9.14-2  |
| Ranger  | org.apache.ranger | unixauthclient                 | 2.4.0.7.1.9.14-2  |
| Ranger  | org.apache.ranger | unixauthservice                | 2.4.0.7.1.9.14-2  |
| Ranger  | org.apache.ranger | unixusersync                   | 2.4.0.7.1.9.14-2  |
| Solr    | org.apache.solr   | solr-analysis-extras           | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-analytics                 | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-cell                      | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-core                      | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-dataimporthandler         | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-dataimporthandler-extras  | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-gcs-repository            | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-jaegertracer-configurator | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-langid                    | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-ltr                       | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-prometheus-exporter       | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-s3-repository             | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-security-util             | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-solrj                     | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-test-framework            | 8.11.2.7.1.9.14-2 |
| Solr    | org.apache.solr   | solr-velocity                  | 8.11.2.7.1.9.14-2 |
| Spark   | org.apache.spark  | spark-avro_2.11                | 2.4.8.7.1.9.14-2  |
| Spark   | org.apache.spark  | spark-catalyst_2.11            | 2.4.8.7.1.9.14-2  |
| Spark   | org.apache.spark  | spark-core_2.11                | 2.4.8.7.1.9.14-2  |
| Spark   | org.apache.spark  | spark-graphx_2.11              | 2.4.8.7.1.9.14-2  |
| Spark   | org.apache.spark  | spark-hadoop-cloud_2.11        | 2.4.8.7.1.9.14-2  |
| Spark   | org.apache.spark  | spark-hive_2.11                | 2.4.8.7.1.9.14-2  |
| Spark   | org.apache.spark  | spark-kubernetes_2.11          | 2.4.8.7.1.9.14-2  |
| Spark   | org.apache.spark  | spark-kvstore_2.11             | 2.4.8.7.1.9.14-2  |
| Spark   | org.apache.spark  | spark-launcher_2.11            | 2.4.8.7.1.9.14-2  |

| Project  | groupId             | artifactId                               | version          |
|----------|---------------------|------------------------------------------|------------------|
| Spark    | org.apache.spark    | spark-mllib-local_2.11                   | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-mllib_2.11                         | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-network-common_2.11                | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-network-shuffle_2.11               | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-network-yarn_2.11                  | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-repl_2.11                          | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-sketch_2.11                        | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-sql-kafka-0-10_2.11                | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-sql_2.11                           | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-streaming-kafka-0-10-assembly_2.11 | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-streaming-kafka-0-10_2.11          | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-streaming_2.11                     | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-tags_2.11                          | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-token-provider-kafka-0-10_2.11     | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-unsafe_2.11                        | 2.4.8.7.1.9.14-2 |
| Spark    | org.apache.spark    | spark-yarn_2.11                          | 2.4.8.7.1.9.14-2 |
| Sqoop    | org.apache.sqoop    | sqoop                                    | 1.4.7.7.1.9.14-2 |
| Sqoop    | org.apache.sqoop    | sqoop-test                               | 1.4.7.7.1.9.14-2 |
| Tez      | org.apache.tez      | hadoop-shim                              | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | hadoop-shim-2.8                          | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-api                                  | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-aux-services                         | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-common                               | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-dag                                  | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-examples                             | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-ext-service-tests                    | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-history-parser                       | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-javadoc-tools                        | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-job-analyzer                         | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-mapreduce                            | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-protobuf-history-plugin              | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-runtime-internals                    | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-runtime-library                      | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-tests                                | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-yarn-timeline-cache-plugin           | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-yarn-timeline-history                | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-yarn-timeline-history-with-acls      | 0.9.1.7.1.9.14-2 |
| Tez      | org.apache.tez      | tez-yarn-timeline-history-with-fs        | 0.9.1.7.1.9.14-2 |
| Zeppelin | org.apache.zeppelin | zeppelin-angular                         | 0.8.2.7.1.9.14-2 |

| Project   | groupId              | artifactId                     | version          |
|-----------|----------------------|--------------------------------|------------------|
| Zeppelin  | org.apache.zeppelin  | zeppelin-display               | 0.8.2.7.1.9.14-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-interpreter           | 0.8.2.7.1.9.14-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-jdbc                  | 0.8.2.7.1.9.14-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-jupyter               | 0.8.2.7.1.9.14-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-livy                  | 0.8.2.7.1.9.14-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-markdown              | 0.8.2.7.1.9.14-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-server                | 0.8.2.7.1.9.14-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-shell                 | 0.8.2.7.1.9.14-2 |
| Zeppelin  | org.apache.zeppelin  | zeppelin-zengine               | 0.8.2.7.1.9.14-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper                      | 3.8.1.7.1.9.14-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-fatjar       | 3.8.1.7.1.9.14-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-loggraph     | 3.8.1.7.1.9.14-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-rest         | 3.8.1.7.1.9.14-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-contrib-zooinspector | 3.8.1.7.1.9.14-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-it                   | 3.8.1.7.1.9.14-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-jute                 | 3.8.1.7.1.9.14-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-prometheus-metrics   | 3.8.1.7.1.9.14-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-election     | 3.8.1.7.1.9.14-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-lock         | 3.8.1.7.1.9.14-2 |
| ZooKeeper | org.apache.zookeeper | zookeeper-recipes-queue        | 3.8.1.7.1.9.14-2 |

## What's new in Platform Support

You must be aware of the platform support for the Cloudera Runtime 7.1.9 release.

### Platform Support Enhancements

New OS Versions: CDP Private Cloud Base now supports the following OS versions

- RHEL 9
- RHEL 8
- RHEL 8.8 FIPS (added RHEL 8.8 support for FIPS customers with JDK8)
- Oracle 8.8 UEK
- Sles 15 SP4

For more information on the support matrix, see [Cloudera Support Matrix](#).

## Fixed issues in Cloudera Runtime 7.1.9

You can review the list of reported issues and their fixes in Cloudera Runtime 7.1.9. Fixed issues represent selected issues that were previously logged through Cloudera Support, but are now addressed in the current Runtime release. These issues may have been reported in previous versions of Runtime as a known issue; meaning they were reported by customers or identified by Cloudera Quality Engineering teams.

### Fixed Issues in Apache Atlas

Review the list of Atlas issues that are resolved in Cloudera Runtime 7.1.9.

**OPSAPS-67782: ZDU | During rolling upgrade one ATLAS server failed to start while other server went fine**

Exclusion of jackson databind from parent dependency

**OPSAPS-67878: [CKP-2,3,4] Apply config injectors for all Atlas csd**

Changes made in 7.1.8 CSD, for 7.1.9 CSD CM team has already taken care.

**CDPD-56309: [Atlas] Download Search with Basic Search gives java.io.FileNotFoundException**

Default value for the download directory will be read from system property "user.dir" instead of property "atlas.home"

**CDPD-49053: Atlas - Upgrade Tinkerpop to 3.5.4**

Upgrade Tinkerpop to 3.5.4

**CDPD-48090: Atlas - Upgrade icu4j to 66.1+ due to CVE-2020-21913**

Upgrade icu4j to 66.1+ due to CVE-2020-21913

**CDPD-50022: Test failure - Atlas 'Service Unavailable' error**

Circular dependency issue caused by repository.src.main.java.org.apache.atlas.tasks.TaskRegistry class is resolved by adding @Lazy annotation on @Component definition of this class.

**CDPD-29307: Kafka producer entity stays in incomplete state in Atlas**

The Kafka-Atlas plugin now fully creates Producer and Consumer entities and won't generate incomplete ones.

**CDPD-57305: Atlas - Upgrade moment.js to 2.29.4 due to CVE-2022-24785, CVE-2022-31129**

Updated moment.js version to 2.29.4

**CDPD-45073: Implement aging for audits stored by Atlas.**

Feature to reduce Atlas audit storage by aging out audit data, using different criteria like TTL and audit count

**CDPD-58592: [CKP4 (same value)] Atlas not null filter on classification returns null values**

Cause: It is because of Solr version upgrade, until 8.4.1, Solr supported non empty string. Fix: For IndexQuery : ["" TO \*] works to get nonEmpty field entities. For Inmemory Predicates: Used NonEmptyPredicate

**CDPD-55098: 7.1.9 - Dynamic Index Recovery issues and improvements**

Introduced AtlasClient API to perform index recovery

**CDPD-50762: Regression : admin/audits , admin/purge fail with "[\_\_AtlasAuditEntry.startTime] is not indexed in the targeted index [vertex\_index]" 7.1.9**

Cause: The attributes of AtlasAuditEntry type where not getting indexed in the Solr, because the entity of AtlasAuditEntry gets created before the attributes gets indexed. Fix: Ordering the typeDefChangeListeners helped, to create the AtlasAuditEntry typeDef first before auditing.

**CDPD-59409: Tags are not getting synced from one node of rangertagsync**

Log4j version incompatibility between Atlas and Ranger led to this issue. A temporary fix has been merged to avoid the exception.

**CDPD-54645: Ranger tag sync for Iceberg table type**

Added iceberg support for Ranger Tagsync

**CDPD-59713: [backport] Indexed string field (solr.StrField) which is too large ERROR**

Impala process entities created by ImpalaHook saves query-string in name field. Since query-string can be large, you are getting the longer than the max error.

To store qualifiedName in name field instead of query-string

**CDPD-49495: OOM issue with DSL search caching added by CDPD-27872**

Fixed OOM seen after large number DSL search requests. This was caused by the fix for ATLAS-4347

**CDPD-57277: Atlas - Upgrade Spring Framework to 5.3.27/6.0.8 due to CVE-2023-20861, CVE-2023-20860 and CVE-2023-20863**

Upgrade Spring Framework to 5.3.27 from 5.3.21

**CDPD-57433: Atlas - Upgrade gremlin shaded to 3.5.5+ due to jackson-databind CVEs**

Upgrade gremlin shaded to 3.5.5 from 3.5.4

**CDPD-49450: Atlas - Upgrade jettison to 1.5.4 due to CVE-2022-45685 and CVE-2022-45693**

Upgrade jettison to 1.5.4 from 1.3.7.

**CDPD-53745: commits in CDH-7.1.8.x but NOT IN CDH-7.1.9.x**

This issue is fix in CDPD-52776.

**CDPD-54846: Atlas: CVE-2023-24998-upgrade commons-fileupload library to version 1.5**

Upgrade commons-fileupload library to version 1.5 from 1.3.3.

**CDPD-54852: Backward compatibility for check provided for AttributeName in Parent and Child TypeDef**

This patch provides backward compatibility for two changes mentioned: <https://issues.apache.org/jira/browse/ATLAS-3872> Restrict typedef creation when a child type attribute conflicts with parent type attribute of same name <https://issues.apache.org/jira/browse/ATLAS-4522>. Updating typedef with new supertype should be allowed only if attributes are unique compared to other existing supertypes.

**CDPD-49451: Atlas - Upgrade snakeyaml due to CVE-2022-1471**

Upgrade snakeyaml to 2.0 from 1.33.

**CDPD-50740: [7.1.9]Atlas - Upgrade azure-storage libraries due to CVE-2022-30187**

Upgrade azure-storage-blob version to 12.9.0 from 12.20.2 Upgrade azure-storage-queue version to 12.7.0 from 12.15.2.

**CDPD-55440: Atlas - Upgrade snakeyaml to 2.0 due to CVE-2022-1471**

Upgrade snakeyaml to 2.0 from 1.33.

**CDPD-58492: Atlas - Upgrade Netty Project to 4.1.94.Final due CVE-2023-34462**

Upgrade Netty Project to 4.1.94.Final from 4.1.86.Final.

**CDPD-49452: Atlas - Upgrade Netty to 4.1.86.Final due to CVE-2022-41881, CVE-2022-41915**

Upgrade Netty to 4.1.86.Final from 4.1.77.Final.

**CDPD-57685: Atlas docs failures in CDH builds**

Added the library "fix-esm" to handle the dependency update issue which was causing the build to fail.

**CDPD-55617: Atlas - Upgrade Nimbus-JOSE-JWT to 9.24 due to CVEs coming from json-smart**

Upgrade Nimbus-JOSE-JWT to 9.24 from 9.8.1.

**CDPD-49978: Atlas - Upgrade icu4j to 66.1+ due to CVE-2020-21913**

Upgrade icu4j to 66.1.

**CDPD-58596: Large number of warn messages logged in tagsync log indicating process entity notification dropped**

Added code to send version in the messages sent from Atlas to Ranger, so that the warning messages are not seen.

**CDPD-53808: Atlas - Upgrade Spring Framework to 5.3.27/6.0.8 due to CVE-2023-20861, CVE-2023-20860 and CVE-2023-20863**

Upgrade Spring Framework to 5.3.27 from 5.3.20.



**CDPD-51181: Atlas - Upgrade Woodstox to 5.4.0/6.4.0 due to multiple CVEs**

Upgrade Woodstox to 5.4.0 from 5.0.3.

**CDPD-49980: Atlas - Upgrade Tinkerpop to 3.5.4**

Upgrade Tinkerpop to 3.5.4 from 3.5.2.

**CDPD-55876: Atlas - Upgrade Spring Security to 5.7.8+/5.8.3+/6.0.3+ due to CVE-2023-20862**

Upgrade Spring Security to 5.8.3 to 5.7.5.

**CDPD-55252: Atlas - Upgrade jackson-databind to 2.12.7.1/2.13.4.1+ due to CVE-2022-42003, CVE-2022-42004**

Upgrade jackson-databind to 2.12.7.1.

**CDPD-50741: [7.1.9]Atlas - Upgrade reactor-netty to 1.0.24+ due to CVE-2022-31684**

Upgrade reactor-netty to 1.0.24.

**CDPD-54864: Every hive insert generates an Atlas audit event**

Introduced "DML audit filters" feature to skip unnecessary DML audit events using configuration property atlas.hook.hive.skip.dml.messages=true (Enabled by default).

**CDPD-48641: Atlas - Support for JDK17 in all sub-components**

Added JDK17 support for Atlas.

**CDPD-56497: Regression : DSL queries redirected to passive server fails**

DSL search request sent to Passive server having HTML encoded characters are now properly redirected to Active server.

**CDPD-59758: [ST][Atlas] test\_export\_import\_api\_sanity test fails**

Upgraded tinkerpops version from 3.5.5 to 3.5.6 : This resolved unsupported class major version error faced while running export API in JDK17 Runtime.

**CDPD-56594: Lineage (spark\_process) is not created for views created on Iceberg tables.**

Added lineage support for creating views to Spark Atlas Connector for Spark 3 which is required for Iceberg tables.

**CDPD-54990: The in-place migration of Hive table to Iceberg table with ALTER TABLE storage\_handler using Beeline creates new iceberg\_table entity but retains the old hive\_table entity as is**

Running the query no longer results with Atlas having two entities with same name but different types. One with hive\_table and another with iceberg\_table.

**Apache patch information**

- ATLAS-4733
- ATLAS-4754
- ATLAS-4442
- ATLAS-4735
- ATLAS-4768
- ATLAS-4762
- ATLAS-4727
- ATLAS-4571
- ATLAS-4576
- ATLAS-4757

**Fixed Issues in Apache Avro**

Review the list of Avro issues that are resolved in Cloudera Runtime 7.1.9.

**CDPD-23451: Avro - Remove/replace jackson-mapper-asl dependency**

Removed jackson-mapper-asl dependency that contains a couple of CVEs from Avro

### Apache patch information

Avro rebased from 1.8.2 to 1.11.1.

## Fixed issues in Cloud Connectors

Review the list of Cloud Connectors issues that are resolved in Cloudera Runtime 7.1.9.

### **CDPD-48449: distcp -update skips files of same size, name when transferring from HDFS to S3**

The Distcp -update option, may encounter potential inaccuracies by skipping the copy when doing incremental update of files with identical names and sizes during the transfer process from HDFS to S3 or ABFS. This occurs due to the absence of checksum verification between the files for different stores. In order to address this concern, we employ the modification time as a means to minimize the occurrence of incorrect skips. If the source file has been modified more recently than its corresponding destination file, we proceed with the copy operation; otherwise, the file is skipped.

### **CDPD-43464: HADOOP-18344 AWS SDK update to 1.12.262 due to CVE-2022-31159**

The aws-java-sdk library was updated to 1.12.262+ due to CVE-2022-31159. Note that the S3A connector has never been vulnerable to the CVE, as it does not use the SDK's TransferManager for downloading files.

### **CDPD-45959: Some tests fail with ssl3\_get\_server\_certificate:certificate verify failed**

"fs.azure.ssl.channel.mode" has been set to "Default\_JSSE". Switch to "Default" if the version of OpenSSL installed in your OS can successfully negotiate SSL connections with Azure to achieve possibly improved performance.

### **CDPD-12425: support S3 client side encryption**

The S3A connector now supports S3-CSE client side encryption. See the documentation for the specific details on how to enable this.

### **CDPD-29477: HADOOP-17618. ABFS: Partially obfuscate SAS object IDs in Logs**

ABFS: Partially obfuscate SAS object IDs in Logs

### **CDPD-35030: HADOOP-18112. Rename operation fails during multi object delete of size more than 1000.**

Fix multi object delete in S3A when number of objects is more than 1000

### **CDPD-46175: HADOOP-18521. ABFS prefetching input stream corruption**

ABFS prefetching input stream corruption

### **CDPD-46543: HADOOP-18526. Leak of S3AInstrumentation instances using Hadoop Metrics references**

Leak of S3AInstrumentation instances using Hadoop Metrics references

### **CDPD-56830: HADOOP-18233. Initialization race condition with TemporaryAWSCredentialsProvider**

Initialization race condition with TemporaryAWSCredentialsProvider

### **CDPD-35182: HADOOP-17198. Support S3 Access Points.**

The S3A connector supports S3 AccessPoints. The access point for a bucket must be set for that specific bucket; for a bucket "NAME" the option would be "fs.s3a.bucket.NAME.accesspoint.arn". If the option "fs.s3a.accesspoint.required" is set to true, then all buckets must be configured with AccessPoint ARNs.

### Apache patch information

- HADOOP-18596
- HADOOP-13887
- HADOOP-17618
- HADOOP-18521
- HADOOP-18233
- HADOOP-18526

## Fixed issues in Cruise Control

Review the list of Cruise Control issues that are resolved in Cloudera Runtime 7.1.9.

### **OPSAPS-68148: Cruise Control rack aware goal upgrade handler**

As Cruise Control automatically overrides all occurrences of the deprecated RackAwareGoal with RackAwareDistributionGoal during upgrade, the customized values of the Cruise Control goals will remain the same and there is no need to manually provide the values after an upgrade.

### **CDPD-47616: Unable to initiate rebalance, number of valid windows (NumValidWindows) is zero**

This issue has been fixed as the Cloudera Manager Metrics Reporter is deprecated from Cruise Control.

### **OPSAPS-66432: Cruise Control does not start with Python 3**

This issue has been fixed and Cruise Control will start successfully with Python 3.

## Fixed Issues in Apache Hadoop

Review the list of Hadoop issues that are resolved in Cloudera Runtime 7.1.9.

### **CDPD-46151: HADOOP-18469: Adding XMLUtils methods to centralize code that creates secure XML parsers.**

This issue is resolved.

### **CDPD-46149: HDFS-16766: Fixing vulnerability related to XML External Entity (XXE) when processing XML received from untrusted sources.**

This issue is resolved.

### **CDPD-35138: HADOOP-17837: Add unresolved endpoint value to UnknownHostException**

HADOOP-17837: Add unresolved endpoint value to UnknownHostException

### **CDPD-35233: HADOOP-17328. LazyPersist Overwrite fails in direct write mode.**

HADOOP-17328. LazyPersist Overwrite fails in direct write mode.

### **CDPD-50248: HADOOP-17125: Simplified SnappyCodec Compression with snappy-java**

The SnappyCodec now uses the snappy-java library, which includes native binaries.

## Apache patch information

- HADOOP-17837
- HADOOP-17328
- HADOOP-17125

## Fixed Issues in Apache HDFS

Review the list of HDFS issues that are resolved in Cloudera Runtime 7.1.9.

### **CDPD-46151: HADOOP-18469: Adding XMLUtils methods to centralize code that creates secure XML parsers.**

This issue is resolved.

### **CDPD-46149: HDFS-16766: Fixing vulnerability related to XML External Entity (XXE) when processing XML received from untrusted sources.**

This issue is resolved.

### **OPSAPS-65324: Logredactor pattern causes failed blockID's logging**

The regex for the social security numbers and the credit card numbers was updated so that it will not redact blockIDs. Added a word boundary `{/b}` for the beginning and the end. This will be default from the CDP 7.1.9 release.

### **OPSAPS-67496: HDFS fails on Java 17 system tests**

Added necessary JVM startup options to be able to run HDFS on Java 17.

**OPSAPS-63558: Previously, DistCp did not correctly report renames and deletes \*in case of snapshot diff based\* HDFS replications. This change extends DistCp's output report to contain counters related to snapshot diff based replications beside the already reported counters. These counters are added to the following group:**

com.cloudera.enterprise.distcp.DistCpSyncCounter. The following, new counters are added:

- FILES\_MOVED\_TO\_COMMON\_TEMP\_DIR: Number of files and directories moved to a common temporary directory to be renamed or deleted later in the process. The common temporary directory referenced here is a sibling of the replication target directory. This counter is the sum of FILES\_DELETED\_VIA\_COMMON\_TEMP\_DIR and FILES\_RENAMED\_VIA\_COMMON\_TEMP\_DIR.
- FILES\_DELETED\_VIA\_COMMON\_TEMP\_DIR: Number of files moved to a common temporary directory to be deleted later. The common temporary directory referenced here is a sibling of the replication target directory.
- FILES\_RENAMED\_VIA\_COMMON\_TEMP\_DIR: Number of files moved to a common temporary directory first, then moved to their final place. The common temporary directory referenced here is a sibling of the replication target directory.
- FILES\_DIRECT\_DELETED: Number of files deleted directly. This is a feature introduced in OPSAPS-63759.
- FILES\_DIRECT\_RENAMED: Number of files renamed directly, without moving to an intermediate temporary directory. This is a feature introduced in OPSAPS-63930.
- FILES\_DIRECT\_RENAMED\_VIA\_TEMP\_LOCATION: Number of files moved to an intermediate temporary directory and then renamed. This intermediate temporary directory is different from the common temporary directory referenced in the FILES\_RENAMED\_VIA\_COMMON\_TEMP\_DIR counter's description. This is also related to OPSAPS-63930. The values of FILES\_DELETED\_VIA\_COMMON\_TEMP\_DIR and FILES\_DIRECT\_DELETED are also aggregated in the replication result as the number of files deleted.

## Fixed Issues in Apache HBase

Review the list of HBase issues that are resolved in Cloudera Runtime 7.1.9.

### CDPD-49477: HBase rebase to 2.4.17

In CDP Private Cloud Base, HBase component is upgraded to base version 2.4.17 and Apache HBase third party to base version 4.1.1.

### Apache patch information

None

## Fixed Issues in Apache Hive

Review the list of Hive issues that are resolved in Cloudera Runtime 7.1.9.

### CDPD-27233: ACID HMS table names should changed to VARCHAR(256)

In the Hive metastore schema, there were some tables where the column width for certain columns was inconsistent across tables. This has been addressed.

### CDPD-34727: Graceful shutdown of HiveServer2

With the introduction of a graceful shutdown state for HiveServer (HS2), HS2 will wait for running (in-flight) queries to complete within a configured amount of time before shutting down.

### CDPD-39708: Use cdpd\_guava\_version from external\_versions.ini

This patch uses guava dependency version from cdpd repro.

### CDPD-43493: Upgrade JUnit to 4.13.2 due to medium CVEs

Upgraded JUnit to 4.13.2 to fix CVEs.

### CDPD-43540: Upgrade jersey to 2.36/3.0.5 due to medium CVEs

Upgraded jersey to 2.36 to fix CVEs.

### CDPD-45017: Backport HIVE-23242: Fix flaky tests testHouseKeepingThreadExistence in TestMetastoreHousekeepingLeaderEmptyConfig and TestMetastoreHousekeepingLeader

Re-enable the HMS leader test in TestMetastoreHousekeepingLeaderEmptyConfig and TestMetastoreHousekeepingLeader.

**CDPD-47132: Pushdown Date data type to Hive metastore via direct sql / JDO**

Fix partition filtering when querying partition metadata from metastore and partition key column data type is date.

**CDPD-47464: Alter view command allowed even when user has a deny policy on the underlying table**

"Alter View As" queries were not authorized correctly. This fix addresses the security concern around the authorization of "Alter View As" queries.

**CDPD-47557: SparklyRHWC certification with R4**

SparklyRHWC supports R version 4.0.5.

**CDPD-48022: Upgrade postgresql to 42.5.1 due to CVE-2022-41946**

Upgraded PostgreSQL to 42.5.1 to fix CVEs

**CDPD-48801: Pushdown Timestamp data type to Hive metastore via direct sql / JDO**

Support partition filtering when querying partition metadata from metastore and partition key column data type is timestamp.

**CDPD-48989: Hive/Impala tests in CDH-7.1.8.x branch fails with DataNucleus connection timeout errors**

This fix introduces a secondary connection pool for HMS DataNucleus value generator to avoid connection starvation.

**CDPD-49145: Oozie and Spark tests are failing in multi-comp-pre with Zookeeper-based or direct JDBC URL to Hive**

HiveServer (HS2) uses InetAddress.getHostName() API to get its hostname and register itself with Zookeeper. The API behaviour is changed on JDK 11 with specific operating systems to return only the hostname without the domain suffix. Therefore, HS2 is inaccessible to clients when the server information is obtained from Zookeeper. To address this, the InetAddress.getCanonicalHostName() API is used to fetch the hostname along with the fully qualified domain name.

**CDPD-49492: Extend batch partition APIs to ignore partition schemas**

This patch addresses issues with high partition workloads. See the [documentation](#) for information about the recommended configurations required to run high partition workloads.

**CDPD-49507: {OWNER} policy not working with HIVE UDFs in RangerHiveAuthorizer**

The UDFs used in Hive will now honor {Owner} policies in Ranger with this fix.

**CDPD-50148: Add double quotes for tables in PartitionProjectionEvaluator**

This fix addresses an issue where a missing relation SQLException is returned when PartitionProjectionEvaluator requests partitions against PostgreSQL.

**CDPD-50450: Backport HIVE-27201: Inconsistency between session Hive and thread-local Hive may cause HS2 deadlock**

Two HiveServer (HS2) sessions can go into a deadlock state due to RANGER-3593 and can indefinitely wait for each other. This fix resolves the deadlock condition.

**CDPD-50464: Installation failures with Hive errors**

The fix sets a different name for the connection pool for Compactor to resolve conflict with the connection metrics of pool for ObjectStore.

**CDPD-50730: Hive WebUI HTTP 500 error due to jar order in classpath**

Removed javax.servlet.jsp-api dependency from HiveServer (HS2) to avoid the intermittent Null Pointer Exception while opening the home web page.

**CDPD-51885: Fix and backport HIVE-27163 - Column stats are not getting published after an insert query into an external table with custom location**

Column stats are published after an insert into the external table created with an empty custom location.

**CDPD-53363: Backport HIVE-25032: Optimize PartitionManagementTask**

Optimize PartitionManagementTask by searching for the required tables to be repaired and dropping the obsolete partitions in a bulk.

**CDPD-55511: CDPD-39232 / RANGER-3593 causing slow execution of SHOW TABLES command**

Currently, TableMeta does not include ownership information which makes it difficult for filterTableMetas to efficiently filter based on \${OWNER} privileges.

Ranger will look up ownership information for database objects (databases/tables/views) through HMS api calls during authorization checks, if it is not provided by the callee. This can be very slow for commands like SHOW TABLES when the database has a large amount of tables, since it may generate a call per object.

This fix addresses the issue by providing ownership information for these commands from the Hive side so that the Ranger authorization plugin is no longer required to make HMS calls in this situation.

**CDPD-55866: Backport HIVE-18827 to CDH-7.1.9.x**

Fixed unnecessary dynamic value exceptions

**CDPD-55867: Backport HIVE-23295 to CDH-7.1.9.x**

This fix addresses a potential Null Pointer Exception when dynamic values are not available while fetching the predicate literal list.

**CDPD-55868: Backport HIVE-23410 to CDH-7.1.9.x**

As part of an earlier fix, insert operation was modified to write directly to the table location instead of the staging directory. The same improvement is now available for ACID update and delete operations.

**CDPD-55869: Backport HIVE-24581 to CDH-7.1.9.x**

The split generation in OrcInputformat is tightly coupled with ACID and AcidUtils.getAcidState is called even when the table is not transactional. The fix includes the following changes:

- Removes unnecessary AcidUtils.getAcidState call from OrcInputformat when the table is not transactional
- Removes redundant filesystem utility functions from AcidUtils to HdfsUtils

Removed AcidUtils call from OrcInputformat for non transactional tables

**CDPD-55870: Backport HIVE-24669 to CDH-7.1.9.x**

Improved FileSystem usage in Hive::loadPartitionInternal to improve performance.

**CDPD-55871: Backport HIVE-24679 to CDH-7.1.9.x**

In order to improve performance, FullDPSpecs in loadDynamicPartitions are reused to avoid double listing.

**CDPD-55873: Backport HIVE-24682 to CDH-7.1.9.x**

The dynamic partition information is collected from the FileSink for direct insert and reused later in the MoveTask.

**CDPD-55875: Backport CDPD-36395 to CDH-7.1.9.x**

Spark/Spark3 builds fail due to Hive changes for CDPD-6264/HIVE-23410

**CDPD-58160: Backport CDPD-56782 to 7.1.9**

This fix addresses an issue where DIRECT\_READER\_V2 mode returned wrong values for string columns after merge query.

**CDPD-59419: Backport CDPD-59091 to 7.1.9**

Return getDefaultSession() if the Spark getActiveSession() returns None.

**OPSAPS-66334: Make modifying hive.security.authorization.sqlstd.confwhitelist.append as first party config**

This fix introduces the `hive.security.authorization.sqlstd.confwhitelist.append` configuration in Cloudera Manager. Users can use this configuration to add Hive properties that are failing instead of the earlier approach of adding this in "Hive Client Advanced Configuration Snippet (Safety Valve) for hive-site.xml".

**OPSAPS-64733: Adding Hive with Ozone warehouse directory fails**

The Hive warehouse directory and external warehouse directory can be set to Ozone during the first Hive run. The Ozone filesystem JAR is available as part of the `hdfs.sh` script (in the `/opt/cloudera/cm/lib/cdh7` directory).

**OPSAPS-67942: Installation failed due to schematool error**

Setting the `hive.hook.proto.base-directory` for Hive Metastore (HMS) in `hive-site.xml` is causing `sys.db` creation to fail because of incompatibility issues between Cloudera Manager 7.11.3 and CDH 7.1.7 SP1/SP2. This patch addresses the issue and sets the above configuration only if the CDH version of Hive is at least CDH 7.1.8.

**OPSAPS-68074: CDH7.1.7SP2 Rolling upgrade to 7.1.9 fails to validate Validating metastore schema tables**

Added upgrade handler for Hive for CDP 7.1.9 release.

**OPSAPS-68213: NFQE JDK17 runtime: module java.base does not issue (java.util.regex, java.io, java.time,java.nio)**

This fix addresses an issue where `java.base` does not issue (`java.util.regex`, `java.io`, `java.time`,`java.nio`).

**Apache patch information**

- HIVE-18827
- HIVE-22193
- HIVE-23242
- HIVE-23295
- HIVE-23410
- HIVE-24581
- HIVE-24669
- HIVE-24679
- HIVE-24682
- HIVE-25032
- HIVE-26049
- HIVE-26419
- HIVE-26640
- HIVE-26701
- HIVE-26778
- HIVE-26787
- HIVE-26850
- HIVE-26893
- HIVE-26914
- HIVE-27091
- HIVE-27116
- HIVE-27147
- HIVE-27163
- HIVE-27179
- HIVE-27201
- HIVE-27285

## Fixed Issues in Hue

Review the list of Hue issues that are resolved in Cloudera Runtime 7.1.9.

### **CDPD-48893: Hue logs get overwritten**

In previous implementations, multiple file handlers wrote to a single log file, causing the Hue logs to be overwritten. Hue now uses a socket handler, which solves this problem.

### **CDPD-47011: Unable to run Hive queries containing the FROM statement in a batch**

Earlier, Hive queries containing the FROM statement would execute and stop when run in a batch. This issue has been resolved by adding the FROM statement to the supported statement types in batch execution.

### **CDPD-44913: Unable to re-upload the same file with changed content in Hue**

Earlier, when you tried to re-upload a file having the same filename and changed contents, the updated file did not get uploaded. This issue has been fixed.

### **CDPD-29285: Deselecting the Enable LDAP TLS option in Cloudera Manager does not work as expected**

Earlier, when you deselected the Enable LDAP TLS option in Cloudera Manager Hue Configuration to enable LDAP authentication using unsecure LDAP (ldap:// instead of ldaps://), the authentication failed with the following error: Caught LDAPError while authenticating ldap: UNAVAILABLE({'info': '00000000: LdapErr: DSID-0C090F77, comment: Error initializing SSL/TLS, data 0, v23f0', 'desc': 'Server is unavailable'}). This issue has been fixed by forcing Hue to use the default value of the use\_start\_tls property (which is false) irrespective of the value present in the ldap\_cert property.

### **CDPD-41497: Unable to upload files to folders on S3/ABFS that contain non-ASCII characters in the folder name**

This issue has been fixed.

### **CDPD-46312: cx\_Oracle 6.4.1 is missing from the CDP stack on 7.1.8 release**

cx\_Oracle 8.3.0 is now included with CDP distribution. You no longer need to download it separately.

### **OPSAPS-64655: Performance issues in loading and using Hue**

You may have experienced a delay in loading the Hue application, running queries, or you saw a blank page while trying to open Hue. This happened because of slower responses from the Hue server due to a limited number of Unicorn worker processes. This issue has been resolved.

### **CDPD-17465: LDAPTest fails with whoami\_s**

Earlier, LDAPTest potentially failed when the extended operation “whoami\_s” was not available. This issue has been resolved.

### **CDPD-43984: Local file import fails with Bad Request(400) error**

Earlier, when you imported an XLSX file in Hue, the import failed with the Bad Request(400) error. This issue has been fixed by increasing the value of the DATA\_UPLOAD\_MAX\_MEMORY\_SIZE property to 5 MB.

## Fixed Issues in Apache Impala

Review the list of Impala issues that are resolved in Cloudera Runtime 7.1.9.

### **OPSAPS-64734: Unable to add Impala as a new service if the default hive warehouse is on Ozone**

Impala creation no longer fails when warehouse directory is on Ozone.

### **OPSAPS-44763: CM should mark an Impala Query if the query touches Erasure Coding file in HDFS/Ozone**

Highlights erasure-coded bytes read in queries and allows filtering on this counter.

### **OPSAPS-65783: Impala Query Monitoring Status Check Alerts**

TypeErrors appeared in Cloudera Manager Agent logs when viewing Impala query monitoring, which caused Impala Query Monitoring Status Check alerts in Cloudera Manager.



Fixed Python3 compatibility issue for Impala query monitoring in Cloudera Manager Agent.

**CDPD-26802: Improved self events detection in catalogd cache**

This parent jira tracks the issues that fixed the consistency issues in detecting self-events in catalogD's cache by using a flag called 'lastSyncEventId' on the metadata object.

**CDPD-47030: Impala-shell ldap\_password\_cmd fails on Python 3.8**

Fixes impala-shell --ldap\_password\_cmd with Python 3.

**CDPD-48721: Impala - Upgrade JQuery Datatables to the latest version to avoid Security issues**

Updates JQuery Datatables in the Impala UI to address CVE-2020-28458 and CVE-2021-23445.

**CDPD-48780: impala-shell now requires setuptools be manually added**

Fixes a regression in CHF3 where impala-shell under Python 2 required installing setuptools.

**CDPD-49015: IMPALA-11859 Metric tracking encrypted bytes read**

Adds BytesReadEncrypted to query profiles and the metric impala-server.io-mgr.encrypted-bytes-read to observe reads of encrypted data from HDFS/Ozone.

**CDPD-50426: Backport IMPALA-11845 to 7.1.9: Fix incorrect check of struct STAR path in resolvePathWithMasking**

Fix IllegalStateException when Ranger column-masking/row-filtering policies are applied on a view and the view alias is used together with STAR in the query, i.e. "v.\*" when "v" is the alias such a view. See more in IMPALA-11845.

**CDPD-49648: Upgrade chart.js to 2.9.4+ due to CVE-2020-7746**

Upgrades chart.js in Impala UI to address CVE-2020-7746.

**CDPD-50186: IMPALA-11966 Enable cache\_ozone\_file\_handles by default**

Enables cache\_ozone\_file\_handles by default to improve scan performance with Ozone.

**CDPD-49781: backport IMPALA-11274 to 7.1.9**

Limits conjunctive normal form (CNF) rewrite of expressions to cases with simple predicates to fix performance regressions.

**CDPD-45661: Support erasure-coding in impala**

Reading erasure-coded files from Ozone is now supported with Impala.

**CDPD-47206: IMPALA-11730 Add support for spilling to Ozone**

Impala can now be configured to spill to Ozone, for example with scratch\_dirs="/tmp/scratch,ofs://ozone-scm:9862/tmp".

**CDPD-8130: Add HTTP Strict Transport Security (HSTS) for Impala**

Adds HTTP Strict Transport Security (HSTS) to Impala UI responses when HTTPS is enabled.

**CDPD-47640: Impala erasure coding support**

Impala now supports interacting with erasure-coded files in HDFS.

**CDPD-47643: Impala SHOW statement to display EC files and policies**

Impala's SHOW FILES, SHOW PARTITIONS, SHOW TABLE STATS, and DESCRIBE EXTENDED now display the erasure code policy for files/tables in filesystems that support erasure coding.

**CDPD-58002: setup-ranger() failed on 7.1.9.x due to RANGER-2895**

This patch does not include any functionality change. It's only related to Impala's own end-to-end test infrastructure.

**CDPD-58314: Impala: run test suite with JDK 17**

Test ran successfully in <https://playground-01.jenkins.cloudera.com/job/impala-private-parameterized/175/>. Automated test runs don't need to hold up the release.

**CDPD-43746: Support for Ozone erasure coded data**

Impala now supports interacting with erasure-coded files in Ozone.

**CDPD-55460: Impala - Upgrade Spring Framework to 5.3.27/6.0.8 due to CVE-2023-20863**

Spring Framework has been upgraded to 5.3.27.

**CDPD-54930: Impala - Support for JDK17 in all sub-components**

Impala supports Java 17.

**CDPD-50912: Unable to connect to impala-shell if there is a file with special character in user's home directory**

impala-shell no longer fails if user's home directory contains special characters.

**CDPD-51180: IMPALA-10186 Write invalid parquet PageLocations which table sort by some columns**

Impala no longer creates empty parquet pages.

**CDPD-41064: IMPALA-11360 Support Java11 in Impala**

Impala supports Java 11.

**CDPD-56557: Rolling Upgrade 7.1.8 to 7.1.9 tests fail when impala-shell gets error 'TException: Invalid response from catalogd for request TGetPartialCatalogObjectRequest'**

This issue was introduced by IMPALA-11350, which added a new field in the middle of Thrift structure TPartialTableInfo and caused backward incompatibility issue between old version of coordinator and new version of catalog server, and new version of coordinator and old version of catalog server. The issue was fixed by checking if the catalog service versions are compatible when registering coordinator and catalog server to a cluster. Incompatible coordinators and catalog server are partitioned in different clusters.

**CDPD-50812: IMPALA-11997 impala-shell: base64.encodestring has been removed in python3.9**

Impala now works with Python 3.9.

**CDPD-56871: Backport fix for IMPALA-12114 to impacted releases**

An issue where idle Impala clients using TLS were needlessly disconnected has been fixed.

**CDPD-57735: IMPALA-12217 cgroup memory limit detection doesn't work for cgroups v2**

Impala now works with cgroups v2.

**Apache patch information**

- IMPALA-11755
- IMPALA-11913
- IMPALA-11892
- IMPALA-11845
- IMPALA-12037
- IMPALA-11274
- IMPALA-7003
- IMPALA-12031
- IMPALA-9487
- IMPALA-11476
- IMPALA-12214
- IMPALA-12013
- IMPALA-12114

**Fixed Issues in Apache Kafka**

Review the list of Kafka issues that are resolved in Cloudera Runtime 7.1.9.

**CDPD-29307: Kafka producer entity stays in incomplete state in Atlas**

The Kafka-Atlas plugin now fully creates producer and consumer entities and does not generate incomplete ones.

**CDPD-48822: AvroConverter ignores default values when converting from Avro to Connect schema**

The AvroConverter now propagates field default values to Connect schemas.

**CDPD-53179: Amazon S3 Sink connector fails when buffer size is reached**

The Amazon S3 Sink connector no longer fails when there is more than 5 MB (buffer size) of data available in a Kafka source topic and the connector receives more than 5 MB of data in a single poll.

**CDPD-45958: Kafka client JAAS override policy validation is incorrect**

The JAAS override filter policy now correctly filters based on the specified rules and does not refuse JAAS configurations because of unknown fields.

**CDPD-44252: Exception during normal operation in MirrorSourceTask causes the task to fail instead of shutting down gracefully**

Stopping the read of offsets in a worker of a MirrorSourceTask will now cause a graceful shutdown and the task can be restarted automatically at a later point.

**OPSAPS-64606: Authorization issues if Kafka Connect is not installed**

If the Kafka Connect role is not present on the cluster, then a Ranger policy (connect internal - topic) is created with default, non-empty topic names. As a result, the Ranger policy include list cannot be empty and will not have any side effects on other Kafka operations.

**OPSAPS-68138: Schema Registry and Kafka cannot download policies from Ranger when using custom Kerberos principals**

Ranger repositories created by the scripts that start Kafka, Kafka Connect, and Schema Registry add both the principal name and the service user name to the repository users and the policies.

**CFM-2966: Stateless NiFi connectors do not work with Java Runtime Environment 9 or later**

Stateless NiFi connectors can now be deployed if you are using Java Runtime Environment (JRE) 9 or later.

**Apache patch information**

- KAFKA-14838

**Fixed Issues in Apache Kudu**

Review the list of Kudu issues that are resolved in Cloudera Runtime 7.1.9.

**SubprocessProtocol has a hard-coded limit on message size, but RangerClient does not honor that while generating requests**

This issue is fixed (KUDU-3450).

**CDPD-59278: Subprocess communication with large messages**

Supports messages of size up to 8MB by default to be transmitted between Kudu master and subprocess server. Given the maximum pipe buffer size is 1MB in linux machines flaky failures are observed if the size of the message is more than 1MB between Kudu master and the subprocess server (KUDU-3489).

**OPSAPS-64525: Kudu can't connect to Ranger KMS with AutoTLS**

If Auto-TLS is enabled on the system, now kudu can use the Auto-TLS certificate to connect to other services. Right now this is only helpful when 'Encrypt data at rest' is enabled so that kudu can connect to Ranger KMS successfully.

**CDPD-47068: Update default value for --tablet\_history\_max\_age\_sec to avoid OOM for kudu-master**

Fixed an issue with kudu-master process consuming too much memory in case of very large clusters, clusters with many thousands of tables, or clusters with huge number of DDL operations per day.

**CDPD-54929: Kudu - Support for JDK17 in all sub-components**

Ran all C++ tests with JDK17 successfully. Built and ran java-example code with JDK17 successfully to test the Java client.

**CDPD-36485: Update replica placement algorithm in kudu-master to avoid range hotspotting**

Kudu's default replica placement algorithm is now range and table aware to prevent hotspotting unlike the old power of two choices algorithm. New replicas from the same range are spread evenly across available tablet servers, the table the range belongs to is used as a tiebreaker.

**Apache patch information**

- KUDU-1945
- KUDU-3248
- KUDU-3326
- KUDU-3406
- KUDU-3413
- KUDU-3418
- KUDU-3437
- KUDU-3448
- KUDU-3450
- KUDU-3451
- KUDU-3452
- KUDU-3455
- KUDU-3472
- KUDU-3474
- KUDU-3475
- KUDU-3476
- KUDU-3479

**Fixed Issues in Apache Knox**

Review the list of Knox issues that are resolved in Cloudera Runtime 7.1.9.

**OPSAPS-67397: Intermittent Knox login error in 7.2.17**

This fix adds CSD support for pac4j.password, which is a pseudo random string that needs to be synced between HA Knox instances for HA SSO to work.

**OPSAPS-67449: Enable Loadbalancing param for Oozie and Impala services in cdp-proxy-api topology**

Sticky session and loadbalancing support was missing for cdp-proxy-api topology, this change adds it back. This change also adds stickysession and LB props for Impala (OPSAPS-67376)

**OPSAPS-63146: Support custom Kerberos path for Knox**

With this change Knox will pick up the krb5 value configured in CM (Administrator -> Settings -> krb5.conf file path ) When the CM property changes, Knox configs will change keeping them in sync.

**OPSAPS-68107: Response code 500 error at large-payload request test over Knox on PC-7.2.17 and DC-7.1.9 SMM executions**

Larger requests (over 15KB) are not failing anymore using the Knox APIs (both SMM UI, and SMM API).

**CDPD-40964: Need to update Knox re-write rules to allow access to newer APIs introduced in Ranger**

Allow metrics,roles, tagrest & xaudit Ranger Admin APIs via Knox proxy

**CDPD-24808: SR with Knox should use round-robin load balancing**

When multiple instances of Schema Registry are running, Knox will use round-robin to forward the requests.

**CDPD-53722: Knox - Upgrade OkHttp to 3.14.9/4.10.0 due to medium CVEs - Pvc**

Upgrade OkHttp to 3.14.9/4.10.0 due to medium CVEs.

**CDPD-50726: [7.1.9.x]- Need to update Knox re-write rules to allow access to newer APIs introduced in Ranger**

Update Knox re-write rules to allow access to newer APIs introduced in Ranger

**CDPD-58562: PvC - Reduce the time taken for Knox startup**

Knox gateway and idbroker startup time improvements were added.

**OPSAPS-58179: HIVE endpoint url is updated on only one Knox host topologies. While on other Knox host, the Cloudera Manager configuration monitoring change is not identified and topologies are not updated with the Hive URL.**

This issue is now fixed.

**CDPD-43069: WEBHDFS operation on Namenode UI via Knox fails when HDFS in HA**

Added failover configuration to WebHDFS to the HaProvider in cdp-proxy topology.

**Apache patch information**

- KNOX-2899
- KNOX-2841

**Fixed Issues in Apache Livy**

Review the list of Livy issues that are resolved in Cloudera Runtime 7.1.9.

**CDPD-55116: Fix Spark vulnerability CVE-2023-22946**

This fix is blacklisting `spark.submit.deployMode` and `spark.submit.proxyUser.allowCustomClasspathInClusterMode` Spark configurations in Livy create session REST API. A new Livy configuration `livy.server.session.allow-custom-classpath` property is added to allow custom class path. If you want to disable or rollback this fix, add `livy.server.session.allow-custom-classpath` as `true` in the Livy configuration using the Cloudera Manager safety valve.

**CDPD-55423: remove verbose output on Livy UI error pages.**

A new `livy.server.send-server-version` Livy configuration property is added. You can set to `true` to send the server version in Cloudera Manager. By default, the value is set to `false`.

**CDPD-48614: Merge latest Apache Livy into CDP 7.1.9**

Livy and Livy for Spark 3 have been updated to upstream version 0.7.2. Additionally, includes some CDP-specific patches and fixes. LDAP is not supported.

**CDPD-45165: Livy HA in CDP PvC Base**

Livy Server Active/passive High Availability is available.

**Apache patch information**

- LIVY-974
- LIVY-975

**Fixed Issues in Navigator Encrypt**

There are no fixed issues for Navigator Encrypt in Cloudera Runtime 7.1.9.

**Fixed Issues in Apache Oozie**

Review the list of Oozie issues that are resolved in Cloudera Runtime 7.1.9.

**CDPD-27164: Oozie should not rely on its LoadBalancer internally**

Oozie will no longer use the LoadBalancer to issue a callback notification, but instead it will try all available Oozie instances one-by-one. If the callback succeeded against one of the Oozie instances, then we will not try the other ones. This way the LoadBalancer will not be used for such purposes.

**CDPD-58538: Oozie should upload and use the config files from sqoop-conf/managers.d when available**

Previously, Oozie did not honor Sqoop's managers.d configurations and extra connector Jars from the lib folder, but now both are automatically available in Oozie's Sqoop action, allowing users to seamlessly utilize connectors like the Sqoop Teradata connector without the need for manual configuration updates or copying Jars to the Workflow's lib folder

**CDPD-50296: Improve Oozie's app state action checking**

Enhanced Oozie's action state checking, to immediately query for running applications right after start-up

**CDPD-41425: LAST\_ONLY and NONE execution modes**

Possible OutOfMemoryError when there are too many coordinator actions to materialize.

If there is a coordinator job defined with a frequency by the minute (e.g. frequency="\* \* \* \* \*"), and start-time lies well in the past, and the coordinator job's execution-mode is LAST\_ONLY or NONE, it can happen that too many CoordinatorActionBean instances are kept on JVM heap within CoordMaterializeTransitionXCommand#insertList as those execution modes omit the check for the throttle value.

As a consequence, we can see as many as multiple hundred thousands of log entries trying to increase CoordMaterializeTransitionXCommand#insertList:

```
[user@host ~]$ grep 'In storeToDB() coord action id' /var/log/oozie/oozie-HOSTNAME.log.out | wc -l478408
```

Apache Jira: <https://issues.apache.org/jira/browse/OOZIE-3254>

**CDPD-43192: Extend Oozie Spark sharelib for HBase interaction**

An additional HBase Jars is added to sharelib to support proper HBase interaction.

**CDPD-43343: Oozie log streaming bug when log timestamps are the same on multiple Oozie servers**

Fixed a bug in the mechanism of the Oozie log streaming.

In case there is a log message in server "A" with the same timestamp as an other log message in server "B", then according to the current implementation, the logs acquired by using `TimestampedMessageParser` corresponding to server "B" will be overwritten by server "A" 's parser (due to the operation of timestampMap.put(earliestParser.getLastTimestamp(), earliestParser)), therefore causing the log messages from server "B" to be ignored from that point.

**CDPD-44209: SqoopMain's printArgs masks Sqoop command line option if preceding one contains "password"**

In Yarn, there was a previous issue in Oozie where command-line arguments were masked incorrectly due to mistaken password detection. As a resolution, customers now have the option to utilize the "oozie.launcher.argumentMaskingExceptionList" configuration. This feature allows them to specify exceptions for password masking. For detailed information on how to use this configuration, please refer to the documentation in oozie-default.xml.

**CDPD-46049: SSH action fails when 'oozie.action.ssh.http.command.post.options' property contains double quotes**

The SSH action's callback mechanism failed with "Invalid content-type" error when capture-output was used in the action definition.

**CDPD-47821: Add missing Sqoop Atlas notification jars to Sqoop share lib**

Earlier, Atlas notification was nonfunctional in Oozie's Sqoop action due to missing Jars, but with the inclusion of those Jars in Oozie's Sqoop ShareLib, Atlas notifications are now expected to function correctly in Oozie's Sqoop action.

**CDPD-56936: Oozie's db cli tool does not honor custom connection properties**

The Oozie DB CLI tool did not respect the "ConnectionProperties" property set by the user through the "oozie.service.JPAService.connection.properties" configuration in Oozie.

**OPSAPS-64457: Make CM provide Oozie the necessary configuration regarding CDPD-43396**

HBase service and Sqoop client dependencies were added for Oozie to have access to their configurations.

**OPSAPS-63816: Configure service hosts to Oozie**

Cloudera Manager will provide the address of all Oozie server instances as a configuration to all Oozie instances. This will be then used by Oozie's callback mechanism so that instead of making the callback through the LoadBalancer in HA mode, the callback will be attempted through each Oozie instance, and if one of them succeeds, then we stop. This way we'll no longer use the LoadBalancer, and make the callback mechanism safer by not having a middle-man.

**OPSAPS-67346: [oozie] Implement validator in CM for Oozie-Spark3 integration**

A validator was added which checks that there is a Spark3 role on all Oozie node. If there is any missing Spark3 role then a warning message will be visible on Oozie's CM page listing the nodes.

**Apache patch information**

- OOOIE-3666
- OOOIE-3254

**Fixed issues in Apache Ozone**

Review the list of Ozone issues that are resolved in Cloudera Runtime 7.1.9.

**SSL Handshake fails between Ozone DataNodes if the two DataNode has their certificate signed by different Ozone Storage Container Managers.**

The issue was fixed due to a change on how these TrustStores are created and configured in Ozone after CDP 7.1.9.

**CDPD-57853: Quota repair count enable quota feature for old bucket/volume.**

This issue is resolved.

**CDPD-49027: Ozone PKI improvements (Cert Rotation, primordial node removal)**

Starting from the CDP 7.1.9 release, the certificates that are there to ensure encrypted communication and authentication between the Ozone internal services are renewed automatically for Ozone Manager, Ozone DataNode, and Ozone Recon Server roles.

The renewal of these certificates by default happens automatically; 28 days before the one year lifetime of these certificates are expiring, without the need of any operator intervention or service disruption.

Storage Container Managers hold a certificate that expires after 5 years from the security bootstrap of the clusters, these certificates still need to be renewed manually after they expire.

**OPSAPS-67940: Snapshot list/restore page**

Snapshot listing page on the OZONE service details where you can see the snapshots of a bucket and delete them or restore them.

**OPSAPS-64733: Adding Hive with Ozone warehouse directory fails**

The Hive warehouse directory and external warehouse directory can be set to Ozone during the first Hive run, as the Ozone filesystem jar is available for the `hdfs.sh` script (it is there in the `/opt/cloudera/cm/lib/cdh7` path)

**OPSAPS-65213: Recommission fails with CSD based services that have decommission support on a host that is in maintenance mode**

In Maintenance host with a commissioned Ozone DataNode or Kafka Broker could not end maintenance, as the command failed at recommissioning either of these nodes.

**OPSAPS-64666: Remove duplicated log directory parameter from Ozone service**

After this change in 7.1.9 or later, only the `log.dir` log directory property will be used for all the logging related files in every Ozone role. For Ozone Prometheus this property was added and will be used.

**OPSAPS-57567: Add rolling restart capability to Ozone Cloudera Manager Service**

Ozone roles can be restarted in rolling fashion from Cloudera Manager.

**OPSAPS-59614: Support SSD profile for Ozone rocksdb**

Support SSD profile for Ozone rocksdb

**OPSAPS-63868: Unable to change Spark log location to Ozone**

The Spark History Location and the Driver Log Location now can be set to an Ozone path, the regex comparison will not fail.

**OPSAPS-57827: Implement Ozone OM and SCM upgrade Needs Finalization canary**

If any Ozone Manager or Storage Container Manager has not been finalized after upgrade, its canary indicator will turn yellow.

**Apache patch information**

- HDDS-8312
- HDDS-7220
- HDDS-9087

**Fixed Issues in Apache Parquet**

Review the list of Parquet issues that are resolved in Cloudera Runtime 7.1.9.

**CDPD-47864: Parquet - CVE-2021-41561-Parquet is vulnerable to Dos attack**

Handle negative values in page headers that fixes CVE-2021-41561

**Apache Patch Information**

- Backport PARQUET-2094
- PARQUET-1633
- PARQUET-1744
- PARQUET-2258
- PARQUET-1964
- PARQUET-1968
- PARQUET-1982
- PARQUET-2117
- PARQUET-2161

**Fixed Issues in Phoenix**

Review the list of Phoenix issues that are resolved in Cloudera Runtime 7.1.9.

**PHOENIX-6881 Implement the applicable Date/Time features from JDBC 4.2**

Phoenix now supports setting and retrieving `java.time.LocalDate`, `java.time.LocalDateTime`, and `java.time.LocalTime` objects through the `ResultSet.getObject()` and `PreparedStatement.setObject()` APIs, as defined by JDBC 4.2.

**Apache patch information****Changes in Phoenix:**

- PHOENIX-6959: Server merges are not used for hinted uncovered indexes for wildcard selects on 5.1
- PHOENIX-6984: Fix fallback to skip-join-merge for hinted global indexes



- PHOENIX-6966: UngroupedAggregateRegionScanner.insertEmptyKeyValue() writes wrong qualifier for encoded CQ tables
- PHOENIX-6965: UngroupedAggregateRegionScanner.insertEmptyKeyValue() generates too many cells
- PHOENIX-6969: While using order in some hinted uncovered queries, a ColumnAlreadyExistsException is thrown
- PHOENIX-6953: Creating indexes on a table with old indexing leads to inconsistent co-processors
- PHOENIX-6873: Use cached connection in IndexHalfStoreFileReaderGenerator
- PHOENIX-6874: Support older HBase versions with broken ShortCircuitConnection
- PHOENIX-6872: Use ServerUtil.ConnectionFactory.getConnection() in UngroupedAggregateRegionScanner
- PHOENIX-6395: Reusing connection instance object instead of creating everytime in PhoenixAccessController class
- PHOENIX-5066: The TimeZone is incorrectly used during writing or reading data
- PHOENIX-6823: Calling Joda-based round() function on temporal PK field causes division by zero error
- PHOENIX-6889: Improve extraction of ENCODED\_QUALIFIERS
- PHOENIX-6720: CREATE TABLE can not recreate column encoded tables that had columns dropped
- PHOENIX-6855: Upgrade from 4.7 to 5+ fails if any of the local indexes exist
- PHOENIX-6818: Remove dependency on the i18n-util library
- PHOENIX-6841: Depend on omid-codahale-metrics
- PHOENIX-6834: Use pooled HConnection for server side upsert select
- PHOENIX-5894: Table versus Table Full Outer join on Salted tables not working
- PHOENIX-6800: Remove superfluous semicolon for import statement in UncoveredLocalIndexRegionScanner
- PHOENIX-6798: Eliminate unnecessary reversed scan for AggregatePlan
- PHOENIX-6653: Add upgrade tests based on HBase snapshots
- PHOENIX-6721: CSV bulkload tool fails with FileNotFoundException if --output points to the S3 location
- PHOENIX-6646: System tables are not upgraded after namespace migration
- PHOENIX-6611: Fix IndexTool -snap option and set VERIFIED in PhoenixIndexImportDirectReducer
- PHOENIX-6601: Fix IndexTools bugs with namespace mapping
- PHOENIX-6480: SchemaExtractionProcessor does not add IMMUTABLE\_STORAGE\_SCHEME and COLUMN\_ENCODED\_BYTES to the generated SQL
- PHOENIX-6509: Forward port PHOENIX-4424 Allow users to create "DEFAULT" and "HBASE" Schema (Uppercase Schema Names)
- PHOENIX-6427: Create sequence fails in lowercase schema
- PHOENIX-6451: Update junit and jcodings versions
- PHOENIX-3067: Phoenix metrics system should not be started in mini-cluster mode
- PHOENIX-6662: Failed to delete rows when PK has one or more DESC column with IN clause
- PHOENIX-6659: RVC with AND clauses return incorrect result
- PHOENIX-6773: PhoenixDatabaseMetadata.getColumns() always returns null COLUMN\_DEF
- PHOENIX-6767: Traversing through all the guideposts to prepare parallel scans is not required for salted tables when the query is point lookup
- PHOENIX-6751: Force using range scan vs skip scan when using the IN operator and large number of RVC elements
- PHOENIX-6771: Allow only "squash and merge" from GitHub UI
- PHOENIX-6766: Fix failure of sqlline due to conflicting jline dependency pulled from Hadoop 3.3
- PHOENIX-6758: During HBase 2 upgrade Phoenix Self healing task fails to create server side connection before reading SYSTEM.TASK
- PHOENIX-6753: Update default HBase 2.4 version to 2.4.13
- PHOENIX-6755: SystemCatalogRegionObserver extends BaseRegionObserver which does not exist in hbase-2.4 branch
- PHOENIX-6733: Ref count leaked test failures
- PHOENIX-6725: ConcurrentMutationException when adding column to table view
- PHOENIX-6530: Fix tenantId generation for Sequential and Uniform load generators

- PHOENIX-5534: Cursors with request metrics enabled throws exception
- PHOENIX-6710: Revert PHOENIX-3842 Turn on back default bloomFilter for Phoenix Tables
- PHOENIX-6705: PagedRegionScanner next throws NPE if pagedFilter is not initialized
- PHOENIX-6699: Phoenix metrics overwriting DefaultMetricsSystem in RegionServers
- PHOENIX-6697: log4j-reload4j is missing from phoenix-assembly
- PHOENIX-6679: PHOENIX-6665 changed column name for CURRENT sequence values
- PHOENIX-6616: Alter table command can be used to set normalization\_enabled=true on salted tables
- PHOENIX-6665: PreparedStatement#getMetaData() no longer fails on parametrized "select next ? values" sequence operations
- PHOENIX-6658: Replace HRegion.get() calls
- PHOENIX-6661: Sqlline does not work on PowerPC linux
- PHOENIX-6636: Replace bundled log4j libraries with reload4j
- PHOENIX-6656: Reindent NonAggregateRegionScannerFactory
- PHOENIX-6645: Remove unnecessary SCN related properties from SYSTEM tables on upgrade
- PHOENIX-6576: Do not use guava's Files.createTempDir()
- PHOENIX-6441: Remove TSOMockModule reference from OmidTransactionProvider
- PHOENIX-6638: Test suite fails with -Dwithout.tephra
- PHOENIX-6591: Update OWASP plugin to latest
- PHOENIX-6579: ACL check does not honor the namespace mapping for mapped views
- PHOENIX-6596: Schema extraction double quotes expressions, resulting in un-executable create statements
- PHOENIX-5865: Column that has default value can not be correctly indexed
- PHOENIX-6615: The Tephra transaction processor cannot be loaded anymore
- PHOENIX-6618: Yetus docker image cannot be built as openjdk 11.0.11 is no longer available
- PHOENIX-6604: Allow using indexes for wildcard topN queries on salted tables
- PHOENIX-6600: Replace deprecated getCall with updated getRpcCall
- PHOENIX-6594: Clean up vararg warnings flagged as errors by Eclipse
- PHOENIX-6592: PhoenixStatsCacheLoader uses non-daemon threads
- PHOENIX-6583: Inserting explicit Null into a (fixed length) binary field is stored as an array of zeroes
- PHOENIX-6577: phoenix\_sandbox.py incompatible with python3
- PHOENIX-6528: When view index pk has a variable length column, read repair does not work correctly
- PHOENIX-6507: DistinctAggregatingResultIterator should keep original tuple order of the AggregatingResultIterator
- PHOENIX-6498: Fix incorrect Correlated Exists Subquery rewrite when Subquery is aggregate
- PHOENIX-6472: In case of region inconsistency Phoenix should stop gracefully
- PHOENIX-6344: CASCADE on ALTER should NOOP when there are no secondary indexes
- PHOENIX-6555: Wait for permissions to sync in permission tests
- PHOENIX-6578: sqlline.py cannot be started from source tree
- PHOENIX-6574: Executing "DROP TABLE" drops all sequences
- PHOENIX-6568: NullPointerException in phoenix-queryserver-client not in phoenix-client-hbase
- PHOENIX-6548: Race condition when triggering index rebuilds as RegionServer closes
- PHOENIX-6558: Update SpotBugs
- PHOENIX-6563: Unable to use 'UPPER'/'LOWER' together with 'IN'
- PHOENIX-6557: Fix code problems flagged by SpotBugs as high priority
- PHOENIX-6556: Log INPUT\_TABLE\_CONDITIONS for MR jobs
- PHOENIX-6550: Upgrade jetty, jackson and commons-io
- PHOENIX-6546: BackwardCompatibilityIT testSystemTaskCreationWithIndexAsyncRebuild is flaky
- PHOENIX-6547: BasePermissionsIT is still a bit flaky
- PHOENIX-6543: De-flake AuditLoggingIT
- PHOENIX-5072: Cursor query loops eternally with local index, returns fine without it
- PHOENIX-6542: WALRecoveryRegionPostOpenIT is flaky

- PHOENIX-6534: Upgrades from previous 4.10 versions are broken
- PHOENIX-6486: Phoenix uses inconsistent chronologies internally, breaking pre-Gregorian date handling
- PHOENIX-6519: Make SchemaTool work with lower case table and column names
- PHOENIX-6518: Implement SHOW CREATE TABLE SQL command
- PHOENIX-6506: Tenant Connection is not able to access or validate Global Sequences
- PHOENIX-6454: Add feature to SchemaTool to get the DDL in specification mode
- PHOENIX-6450: Checkstyle check is creating warnings for line length which are greater than 80
- PHOENIX-6413: Having cannot resolve alias
- PHOENIX-6476: Index tool when verifying from index to data does not correctly split page into tasks
- PHOENIX-6515: Phoenix uses hbase-testing-util but does not list it as a dependency
- PHOENIX-6405: Stop bulk loading into non-empty tables with global secondary indexes
- PHOENIX-6510: Double-checked locking field must be volatile
- PHOENIX-6514: Exception should be thrown

#### Changes in phoenix-connectors:

- PHOENIX-6683: Surround the OR filters with parentheses while converting Spark filters to Phoenix expressions
- PHOENIX-6694: Avoid unnecessary calls of fetching table meta data to region servers holding the system tables in batch oriented jobs in spark or hive otherwise those RS become hotspot
- PHOENIX-6590: Handle rollbacks in Phoenix Spark connector and add way to control batch wise or task wise transactions
- PHOENIX-6566: Shaded Phoenix connectors include restrictive log4j config files
- PHOENIX-6524: Hive connector returns empty AND expression when all children are pushed down
- PHOENIX-6490: Fix shaded Phoenix Hive connector jar name and reduce the size

#### Changes in phoenix-omid:

- OMID-245: Add dependency management for Guava to use 32.1.1
- OMID-244: Upgrade SnakeYaml version to 2.0
- OMID-241: Add logging to TSO server crash
- OMID-242: Bump Google Guice version to 5.1.0 to support JDK 17
- OMID-239: OMID TLS support
- OMID-214: Upgrade commons-io to 2.11.0
- OMID-208: Pass additional options to omid.sh
- OMID-202: Refactor Omid to use Netty 4
- OMID-191: Fix missing executable permission because of MASSEMBLY-941
- OMID-200: Omid client cannot use kerberos cache when using proxyUser
- OMID-199: Omid client cannot use pre-authenticated UserGroupInformation.getCurrentUser()
- OMID-194: OmidTableManager cannot create commit and timestamp tables in kerberos cluster
- OMID-192: Fix missing jcommander dependency

#### Changes in phoenix-queryserver:

- PHOENIX-6908: KerberosName\$NoMatchingRule exception in QueryServer.PhoenixRemoteUserExtractor
- PHOENIX-6704: sqlline-thin.py does not work with python3
- PHOENIX-6762: Phoenix QueryServer cannot run correctly with python 3.8+
- PHOENIX-6727: get\_view\_names() returning empty list
- PHOENIX-6654: queryserver.py sets the umask to 0000 when starting PQS in daemon mode

### Fixed Issues in Apache Ranger

Review the list of Ranger issues that are resolved in Cloudera Runtime 7.1.9.

#### **CDPD-53435: [7.1.9.x] Add/ Update metric details for Ranger TagSync**

Add Metrics APIs for Ranger Tagsync.

#### **CDPD-44451: Add/ Update metric details for Ranger UserSync**

Add Metrics APIs for Ranger Usersync.

**CDPD-58506: User is not allowed delete directory in ozone even though user has permissions**

User is not allowed delete directory in ozone even though user has permissions

**CDPD-69608: Ranger TagSync is out of memory**

Ranger TagSync process runs out of heap memory. The out of memory issue could be because the Kafka messages that are of no interest to TagSync are stored in TagSync process and shall not be cleared until a Kafka message that must be uploaded to Ranger admin is encountered. Leads to build up of messages in a Kafka queue ("ATLAS\_ENTITIES") under the same condition.

**CDPD-50662: [7.1.9.x] - Groups are not visible in mask and row level policy listing tables.**

Groups listing are not visible in mask and row-level policy listing tables.

**CDPD-51892: CLONE - Tag-based policy UI to not show permissions in deny/exception for services that don't support deny/exception**

tag-based policy UI should not show permissions in deny and exception policy-items for service-types that don't support deny and exceptions i.e., service-defs having options.enableDenyAndExceptionsInPolicies=false.

**CDPD-55048: KafkaAuthorization ACL operation Interface implementation in RangerKafkaAuthorizer**

KafkaAuthorization ACL operation Interface implementation in RangerKafkaAuthorizer

**CDPD-57073: RangerClient#createRole singletonMap causes Ozone tenant creation failure in custom-kerberos-principal-option4**

Reverted a change(part of another review for JWT changes) to fix tenant creation in Ozone. Fixes the REST API call, passing auth\_type as kerberos in the request was the issue.

**CDPD-49182: [7.1.9] Ranger AD User Sync - support for AD group names containing slashes**

Adds support for LDAP user and group names with special characters.

**CDPD-44902: Ranger admin feature to delete all external users**

Introduced new feature with the addition of 2 new REST APIs to force delete external users at scale.

**CDPD-46248: Ranger RMS Field issues**

Fixed issues listed in the description below. Please ensure that before applying the patch, RMS service is stopped and the existing RMS resource-mapping is cleaned up. This can be achieved by updating the RMS database tables with the following SQL commands. delete from x\_rms\_resource\_mapping; delete from x\_rms\_service\_resource; delete from x\_rms\_notification; update x\_rms\_mapping\_provider set last\_known\_version=-1; After applying the patch and restarting RMS server, the resource-mappings will be re-synced from HiveMetaStore.

**CDPD-50668: CLONE 7.1.9 - HA support for Ranger User Sync**

This is a new feature which enables support for Ranger usersync in HA(Active-Passive) mode.

**CDPD-48978: kms get currentversion api is returning old keymaterial after key migration from KTS to KMS**

while exporting keys for KTS migration, key version should be in opposite order

**CDPD-49334: Key migration from KTS to RangerKMS**

Key migration from KTS to RangerKMS DB

**CDPD-55419: Ranger - Upgrade json-smart to 2.4.10 due to CVE-2023-1370**

Upgrade json-smart to 2.4.10

**CDPD-53858: metrics are not getting dumped in /var/log/ranger/kms/ranger\_kms\_metric.log file when KMS is stopped**

After discussion internally, it was agreed to dump the metric state in the same regular kms log file when service goes down.

**CDPD-57318: Ranger - Upgrade jackson-dataformat-xml to 2.13.5 due to multiple CVEs in woodstox**

Use woodstox-core to 5.4.0 version

**CDPD-56463: [7.1.9] - Ranger - Upgrade Spring Security to 5.7.8+/5.8.3+/6.0.3+ due to CVE-2023-20862**

Upgrade Spring Security to 5.7.8

**CDPD-50537: [7.1.9.x] - Ranger - Upgrade Kerby to 2.0.3 due to CVE-2023-25613**

Upgrade Kerby to 2.0.3

**CDPD-55561: Ranger - Upgrade bcpkix-jdk15on to 1.70+ due to CVE-2019-17359**

Upgrade bcpkix-jdk15on to 1.70

**CDPD-15744: HA support for Ranger Tag Sync/User Sync**

HA support for Ranger TagSync and UserSync added as part of this new feature enhancement.

**CDPD-54854: CLONE [7.1.9]- Ranger audit metrics deletion is failing**

Code fix for Ranger audit metrics deletion failing.

**CDPD-50648: CLONE [7.1.9] - Ranger is opening a lot of zk connections when solr is down**

Making sure that Ranger closes the Zookeeper connection in case when Solr service is not reachable. Also following the configured number of retries to connect to Solr and on given time intervals.

**CDPD-49503: [Ranger UI] [React JS] If the url to edit a policy, service or permissions for a module, and the url to view user/group/roles contains an invalid id, then page should display an error**

1) If the user enters the Wrong URL in ranger UI It will give 404 Page not found Error page. 2) If the user enters the wrong ID that is not present in the database It will show 400 Data not found page

**CDPD-54619: [7.1.9.x]- Regression caused by CDPD-45891**

Fix uri for getDeletedGroups() in PolicyMgrUserGroupBuilder

**CDPD-44227: Ranger improvement - Roles Import/export API for ranger admin**

Add Roles Import/export API for ranger admin

**CDPD-44198: shell script to export, transform, import of ranger tags for ranger replication**

shell script to export, transform, import of ranger tags for ranger replication

**CDPD-50457: [7.1.9 CLONE] - Provide option to update group memberships when same users/groups are synced from different sync sources**

Update group memberships when same users/groups are synced from different sync sources

**CDPD-56737: Ranger - Upgrade Tomcat to 8.5.89 due to CVE-2023-28709**

Upgrade Tomcat to 8.5.89

**CDPD-50454: [7.1.9.x]- Unable to delete the user if policy is created by same user and added in the policy item**

Allow delete user operation if policy is created by same user and added in the policy item

**CDPD-56300: Introduce config within Ranger to control retention period of x\_auth\_session data**

Add config within Ranger to control retention period of x\_auth\_session table data

**CDPD-55459: Ranger - Upgrade Spring Framework to 5.3.27/6.0.8 due to CVE-2023-20863**

Upgrade Spring Framework to 5.3.27

**CDPD-49638: [7.1.9.x] - Log4j2 support in Ranger**

Added Log4j2 support in Ranger

**CDPD-11878: Support for avoiding multiple access request enrichment**

Optimization to enrich the request only once to alleviate the performance overhead.

**CDPD-50533: [7.1.9.x] - Add unique constraint on resource\_signature column of x\_rms\_service\_resource table**

Add unique constraint on resource\_signature column of x\_rms\_service\_resource table

**CDPD-50605: ArrayIndexOutOfBoundsException exception may be thrown while processing events**

Fix to handle ArrayIndexOutOfBoundsException exception while processing events

**CDPD-49650: [7.1.9.x] - Add Oracle SSL support in ranger**

Oracle SSL Connection support in ranger

**CDPD-58569: Ranger - Upgrade Guava to 32.0.1 due to CVE-2023-2976**

Upgrade Guava library version to 32.0.1

**CDPD-52749: [7.1.9.x]- [Ranger][UserSync]Enumerate Groups will give error on executing 'getent group' command**

Fix for Enumerate Groups will give error on executing 'getent group' command

**CDPD-50368: [7.1.9]- Ranger - Upgrade snakeyaml due to CVE-2022-1471**

Upgrade snakeyaml to 2.0

**CDPD-50433: [7.1.9.x] - No policy found for given version in Ranger Audit page**

Record policy data history during ranger upgrade

**CDPD-49704: deleteUserGroupUtil.py fails to delete username with space**

Allow deletion of users having space in username

**CDPD-58493: Ranger - Upgrade Netty Project to 4.1.94.Final due CVE-2023-34462**

Upgrade Netty Project to 4.1.94.Final

**CDPD-56457: [7.1.9] - Ranger - Upgrade Nimbus-JOSE-JWT to 9.24 due to CVEs coming from json-smart**

Upgrade Nimbus-JOSE-JWT to 9.31

**CDPD-40385: Ranger RMS for Ozone**

This is a new feature introduced in CDP 7.1.9. Ranger RMS will support authorization for Ozone storage locations. RMS for Ozone will co-exist with Hive-HDFS ACL sync and provide authorization for both HDFS and Ozone file systems.

**CDPD-53830: [7.1.9.x] Add/ Update metric details for Ranger RMS**

Add Metrics APIs for Ranger RMS

**CDPD-50564: Add/ Update Additional metric details for Ranger RMS**

Add Additional Metrics for Ranger RMS.

**CDPD-55050: Support SELF\_OR\_PREFIX resource matching scope in Ranger Authorization**

API to find whether a user/group/role is authorized to the given operation on any resource of give type

**CDPD-50670: CLONE 7.1.9 - HA support for Ranger TagSync**

This is a new feature which enables support for Ranger TagSync in HA(Active-Passive) mode.

**CDPD-35034: [SDX/SaaS Migration] Utilities to migrate Ranger Service Tags**

Utilities to migrate Ranger Service Tags

**CDPD-47989: Ranger - Upgrade Netty to 4.1.86.Final due to CVE-2022-41881, CVE-2022-41915**

Upgrade Netty to 4.1.86.Final

**CDPD-49711: assignPermissionToUser in XUserMgr creates entries with NULL moduleId in x\_user\_module\_perm**

Fixed assignPermissionToUser in XUserMgr to correct the bug which assigns permissions for a module (which does not exist) to users with Auditor role.

**CDPD-39208: Review and remove unused RDBMS tables used by Ranger admin service**

Remove unused RDBMS tables used by Ranger admin service

**CDPD-53805: Ozone\_key tag based policies are not working**

What was the Root Cause? Ozone qualified name parsing had a issue wherein '/' was getting included in the key name which resulted in wrong key matching while enforcing policy How was this Issue Resolved? Logic for parsing ozone qualified name changed such that '/' is not included in the key name which was causing issue previously.

**CDPD-55572: shell script to export, transform, import of ranger Roles for ranger replication**

Shell script to export, transform, import of ranger Roles for ranger replication

**CDPD-43132: Allow roles, tagrest & xaudit Ranger Admin APIs via Knox proxy**

This fix allows access to ranger role, tagrest and xaudit ranger admin APIs from Knox proxy.

**CDPD-57018: Ranger - Upgrade aws-java-sdk to 1.12.367+**

Upgrade aws-java-sdk to 1.12.481

**CDPD-48119: Ranger - Upgrade OWASP Java HTML Sanitizer due to security CVEs**

Upgrade OWASP Java HTML Sanitizer

**CDPD-50588: [719 CLONE] - Update dependencies to support macOS aarch64 M1 (Apple Silicon) environment**

Support ranger build on macOS aarch64 M1 (Apple Silicon) environment

**OPSAPS-67025: CM changes for Key migration from KTS to RangerKMS**

Migrating Hadoop keys from Ranger KMS KTS database to Ranger KMS database

**OPSAPS-67374: [7.1.9.x] Unable to locate appender KMS-AUDIT & KMS-METRICS error shown during Ranger KMS start task**

Resolved log4j2 appender issues for Ranger KMS.

**OPSAPS-65704: Alert or notification has to be done when Solr is down resulting in audit pile up in spool directory**

A health alert will be shown on Cloudera Manager for Ranger plugin supported services, when the used space of ranger plugin spool directory (local directory) is greater than the threshold value.

**OPSAPS-65894: Support LunaClient 10.3 for Ranger KMS DB**

New doc created for the Luna 10.5

**CDPD-50726: [7.1.9.x]- Need to update Knox re-write rules to allow access to newer APIs introduced in Ranger**

Update Knox re-write rules to allow access to newer APIs introduced in Ranger

**CDPD-29102: Ranger - Remove log4j 1.x dependencies due to EOL**

Log4j 1.x dependency is removed and upgraded to log4j2

**CDPD-54698: [7.1.9.x] - Ranger - Upgrade Scala to 2.13.9 due to CVE-2022-36944**

Upgrade Scala to 2.13.9 as part of CVE fix

**CDPD-55164: ranger policy replication transform step is not printing logs**

Improve ranger policy replication transformation logs

**CDPD-56462: [7.1.9] - Ranger - Upgrade BeanShell to 2.1b5 due to high CVEs**

Upgrade BeanShell to 2.1b5 by upgrading testNG to 7.0.0

**CDPD-56454: [7.1.9]- Ranger - Upgrade Apache Derby due to critical CVEs**

Upgrade Apache Derby to 10.14.2.0

**CDPD-56455: [7.1.9] - Ranger - Upgrade Spring LDAP to 2.4.1 due to high CVEs**

Upgrade Spring LDAP to 2.4.1

**CDPD-55920: Turning usersync debug logging on results in users not getting synced due to NPE**

Fix NPE while logging debug messages

**CDPD-46256: Ranger Audit metrics page broken in New UI**

Fixed Audit metrics not loading in New UI

**CDPD-48041: Ranger - Upgrade commons-net to 3.9.0 due to CVE-2021-37533**

Upgrade commons-net to 3.9.0

**CDPD-47900: Log4j2 support in Ranger**

Log4j 1.x dependency is removed and upgraded to log4j2

**CDPD-46233: Knox plugin is not working**

Knox service was failing when Audit metrics was enabled. Fix was done to handle the CNF error in Knox Ranger plugin which took care of this error

**CDPD-53826: Ranger - Upgrade jettison to 1.5.4 due to CVE-2023-1436**

Upgrade jettison to 1.5.4

**CDPD-53804: Ranger - Upgrade Spring Framework to 5.3.26/6.0.7 due to CVE-2023-20861 and CVE-2023-20860**

Upgrade Spring Framework to 5.3.27

**CDPD-48032: Ranger - Upgrade jettison to 1.5.2 due to CVE-2022-45685 and CVE-2022-45693**

Upgrade jettison to 1.5.2

**Apache patch information**

- RANGER-4163
- RANGER-4163
- RANGER-4205
- RANGER-3957
- RANGER-4245
- RANGER-4245
- RANGER-4245
- RANGER-3498
- RANGER-3975
- RANGER-4074
- RANGER-4108
- RANGER-4173
- KNOX-2911
- RANGER-4262
- RANGER-3863
- RANGER-4212
- RANGER-4232
- RANGER-4204
- RANGER-4159
- RANGER-3947
- RANGER-4081
- RANGER-4135
- RANGER-4026
- RANGER-4257
- RANGER-4127
- RANGER-4255
- RANGER-4220
- RANGER-4109
- RANGER-4129
- RANGER-4043
- RANGER-4123



- RANGER-3794
- RANGER-4226
- RANGER-4165
- RANGER-4151
- RANGER-4150
- RANGER-4230
- RANGER-4071
- RANGER-3939
- RANGER-4083
- RANGER-4073

## Fixed Issues in Schema Registry

Review the list of Schema Registry issues that are resolved in Cloudera Runtime 7.1.9.

**OPSAPS-66356: Schema Registry's integration with Atlas does not work in secure clusters where Ranger authorization is enabled**

Atlas integration with Schema Registry works out of the box on new clusters.

**CDPD-54379: KafkaJsonSerializer and KafkaJsonDeserializer do not allow null values**

The KafkaJsonSerializer and KafkaJsonDeserializer now properly translates null payloads as null.

**CDPD-49217 and CDPD-50309: Schema Registry caches user group membership indefinitely**

Kerberos users and user groups are not cached in Schema Registry anymore. Any changes to user and user group authentication will take effect without restarting the Schema Registry service.

**CDPD-56890: New schemas cannot be created following an upgrade**

Schemas can be created again after an upgrade even if the latest version of the schema was deleted before the upgrade.

**CDPD-48568: JAR storage does not work on AWS S3 for Schema Registry**

Schema Registry Amazon S3 JAR storage now functions correctly.

**CDPD-48822: AvroConverter ignores default values when converting from Avro to Connect schema**

The AvroConverter now propagates field default values to Connect schemas.

**CDPD-48888: Schema Registry generates redundant schemas when byte[] with default field exists**

Schema Registry's schema normalization and fingerprinting mechanism has been enhanced to properly handle default values for bytes data types.

**CDPD-20977: Add RAW Avro JSON Schema API for Hive Integration**

Added new endpoints where the client can GET the actual schema text as a JSON document. These endpoints were added as a subresource (.../schemaText) to the existing schema version resource endpoints, corresponding to the schemaText property of those, for example, /api/v1/schemaregistry/schemas/{name}/versions/latest/schemaText.

**CDPD-58265: Schema Registry Client incorrectly applies SSL configuration**

The Cloudera distributed Schema Registry Java client applies the SSL configurations correctly even with concurrent access in Jersey clients.

**CDPD-49470: Schema Registry Client retries requests more than the configured maxAttempts when multiple URLs are used**

The Cloudera distributed Schema Registry Java client handles each request as one attempt, and does not attempt more retries based on the number of Schema Registry server URLs anymore.

**OPSAPS-68139: Schema Registry does not apply cluster wide Kerberos principal mapping by default**

The Schema Registry Kerberos Name Rules property is now empty by default. Schema Registry now automatically applies the cluster-wide auth-to-local (ATL) rules by default. During an upgrade, the previously configured value is preserved. If you have been using the default or a custom value, you must manually clear the property following an upgrade to transition to the new default value.

**OPSAPS-68171: Schema Registry does not set Knox principal and service user as trusted proxies when using custom Kerberos principals**

Schema Registry now automatically sets both the Knox principal name and service user name as trusted proxy users.

**OPSAPS-68138: Schema Registry and Kafka cannot download policies from Ranger when using custom Kerberos principals**

Ranger repositories created by the scripts that start Kafka, Kafka Connect, and Schema Registry add both the principal name and the service user name to the repository users and the policies.

**CDPD-55381: Schema Registry issues authentication cookie for the authorized user, not for the authenticated one**

Schema Registry authentication cookie contains the correct authenticated user, even if the authenticated and the authorized users are different. Authenticated and authorized users can be different in scenarios where Schema Registry is used behind Knox.

**CDPD-48853: Schemas created with the Confluent Schema Registry API cannot be viewed in the UI**

Schemas created in Cloudera Schema Registry using the Confluent Schema Registry API are now visible in the Cloudera Schema Registry UI.

In addition, the `/api/v1/schemaregistry/search/schemas/aggregated` endpoint of the Cloudera Schema Registry API now correctly returns schemas created with the Confluent Schema Registry API.

## Fixed Issues in Apache Solr

Review the list of Apache Solr issues that are resolved in Cloudera Runtime 7.1.9.

**CDPD-52804: Fixes Python 3 compatibility issues in HBase-indexer**

HBase indexer failed to start when running with Cloudera Manager built with Python 3 support.

## 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 CDP updates for the public cloud and private cloud base.



**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

Review the list of Spark issues that are resolved in Cloudera Runtime 7.1.9.

### **CDPD-42599: Spark - Update log4j1 to reload4j**

Migrated log4j1 to reload4j to avoid CVE

### **CDPD-58080: Backport SPARK-32951 to Spark 2**

Foldables can be propagated from the Aggregate function.

### **CDPD-50679: Backport CDPD-47129 to 7.1.9**

Now handles empty CSV fields using OpenCSVSerde.

### **CDPD-50203: Backport SPARK-27254 to 7.1.9**

The cleanup completes but becomes invalid in output files for ManifestFileCommitProtocol if the job is aborted.

### **CDPD-50205: Backport SPARK-32638 to 7.1.9**

Prefiously, the WidenSetOperationTypes in a subquery attribute was missing.

### **CDPD-50206: Backport CDPD-43553 to 7.1.9**

Jersey was upgraded to 2.36 to avoid common vulnerabilities and exposures (CVE).

### **CDPD-50161: Backport CDPD-47449 to 7.1.9**

Previously, Spark job failed with NPE while adding kafka-log4j-appender to the classpath.

### **CDPD-50202: Backport SPARK-27210 to 7.1.9**

This patch proposes ManifestFileCommitProtocol to clean up incomplete output files in task level if task aborts.

### **CDPD-52721: Sqoop - Replace log4j 1.x with reload4j**

The log4j was replaced with reload4j in Sqoop.

### **CDPD-43434: Implement support for preventing incompatible log4j classes to be loaded in Sqoop**

A safe-guard was put in place to ensure Sqoop always loads the correct logging related Jars independently from the classpath order.

## Apache patch information

- SPARK-27210
- SPARK-32638
- SPARK-27254
- SPARK-32951

## Fixed Issues in Apache Sqoop

Review the list of Sqoop issues that are resolved in Cloudera Runtime 7.1.9.

### **CDPD-44397: Implement ORC support in Sqoop-Connector-Teradata component**

A new version of Cloudera Connector Powered by Teradata version 1.8.5.1c7 is released which includes ORC support in the Sqoop-Connector-Teradata component. You can use Teradata Manager to import data from the Teradata server to Hive in ORC format.

**CDPD-44431: Disable the Sqoop direct mode feature with ability to enable it again temporarily**

Sqoop's direct mode is no longer supported and is disabled by default. However, you can still 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-44531: Sqoop cannot export Parquet data due to ClassCastException**

Sqoop can now export the following data types from Avro and Parquet files:

- Int, Float, Double to the same RDBMS types
- Long to BigDecimal, Date, Time, TimeStamp
- Bytes to BigDecimal
- Fixed to Decimal and TimeStamp

Note that Fixed to TimeStamp does not work if the source date is based on the Julian calendar.

**CDPD-47175: Sqoop Hive import with ORC file fails with ClassCastException**

The import process of Sqoop to ORC file has been updated. Whenever an unsupported conversion is attempted, Sqoop now provides a comprehensive error message describing the issue.

Sqoop can now import the following data types:

- Byte, Short, Int, Long, Float, Double from the same RDBMS types
- BigDecimal to Long, Double, String
- Date, Timestamp to String, Date, Timestamp

**CDPD-50423: Sqoop ClassCastExceptions when exporting from Parquet**

Sqoop has been enhanced to support additional data type mappings when exporting from Parquet.

**CDPD-56523: Sqoop does not take --hive-compute-stats option into account for hs2-url Hive imports**

Sqoop now considers the `--hive-compute-stats` option for Hive imports when `hs2-url` parameter is used.

## Fixed Issues in Streams Replication Manager

Review the list of Streams Replication Manager issues that are resolved in Cloudera Runtime 7.1.9.

**OPSAPS-67772: SRM Service metrics processing fails when the noexec option is enabled for /tmp**

The SRM Service Kafka Streams application now uses the Kafka Streams state directory to extract the RocksDB .so files

**OPSAPS-67738: SRM Service role's Remote Querying feature does not work when the noexec option is enabled for /tmp**

The SRM service does not add Netty native libraries to /tmp by default as `streams.replication.manager.service.netty.native.working.dir` configuration was introduced.

**OPSAPS-67742: The SRM Service role fails to start if properties are added to Additional Configs For Streams Application Running Inside SRM Service**

The SRM Service role no longer fails to start if properties are added to the **Additional Configs For Streams Application Running Inside SRM Service** configuration. It is also possible to configure the internal Kafka Streams application of the SRM Service role.

**CDPD-60426: Configuration changes are lost following a rolling restart of the service**

SRM no longer fails to apply configuration changes if it is restarted with a rolling restart.

## Fixed Issues in Streams Messaging Manager

Review the list of Streams Messaging Manager (SMM) issues that are resolved in Cloudera Runtime 7.1.9.

**OPSAPS-68158: SMM does not apply cluster wide Kerberos principal mapping by default**

The Kerberos Name Rules property is now empty by default. SMM now automatically applies the cluster-wide auth-to-local (ATL) rules by default. During an upgrade, the previously configured value is preserved. If you have been using the default or a custom value, you must manually clear the property following an upgrade to transition to the new default value.

**OPSAPS-68172: SMM does not set Knox principal and service user as trusted proxies when using custom Kerberos principals**

SMM now automatically sets both the Knox principal name and service user name as trusted proxy users.

**OPSAPS-67575: SMM's Schema Registry client might fail Ranger authorization if mTLS is enabled for Schema Registry**

The Schema Registry client used by SMM no longer includes keystore properties for mTLS when Kerberos is enabled. As a result, even if mTLS is enabled for the Schema Registry server, the Kerberos principal is used for authentication and authorization with Ranger. This fixes possible authorization failures.

**OPSAPS-68107: Response code 500 error when large-payload requests are sent over Knox**

SMM requests that are over 15 KB in size no longer fail when they are sent using the Knox APIs. This fix applies to both SMM UI and SMM REST API requests.

**CDPD-46728: SMM UI shows the consumerGroup instead of the instances on the Profile page's right hand side**

The **Consumer Group Profile** page now correctly shows the consumer instances on the right hand side. Previously the consumer groups were shown.

**CDPD-46465: Searching for workers on the connector overview page freezes the page**

Using the search field on the **Connect Cluster Profile** tab no longer freezes the page.

**CDPD-45406: The Connector Profile page of unassigned connectors is blank**

The **Connector Profile** page of unassigned connectors are now correctly rendered and display that the connector is in an unassigned status.

**CDPD-46073: Data Explorer loads indefinitely**

The **Data Explorer** page no longer breaks if the partition parameter is manually removed from the URL.

**CDPD-49227: The Cluster Replications page crashes if the co-located cluster unknown to SRM**

The **Cluster Replications** page is now correctly displayed even when the co-located Kafka cluster is unknown to SRM.

**CDPD-56086: The Data Explorer modal displays the messages of the wrong topic**

The **Data Explorer** modal that you open by clicking  on the **Topics** page now displays the messages of the selected topic.

**CDPD-49696: Certain alerts may crash the Alerts page**

Composite alerts with one of the conditions containing an assertion on cluster metrics no longer crashes the UI.

**CDPD-54703: Topic Details page does not display the lag of the consumer group**

The Topic Details page no longer incorrectly displays 0 as lag for the consumer groups.

**CDPD-33699: Remove "adjustTopicOverviewMetrics" from SMM**

Removed the logic introduced in 7.1.4/7.2.2.0, where in case the topicMetrics (bytes in/bytes out/messages in) are smaller for a larger time period, the smaller timeperiod's metrics will be displayed. For example, if the metrics are smaller for 30 days and then for 7 days, the 7 day metrics would be used.

**CDPD-43387: Broker Details page does not show the Cloudera Manager button**

The Cloudera Manager buttons that navigate to the broker resource within Cloudera Manager were not visible in previous releases. Now you can navigate from the SMM's broker view to the Cloudera Manager's broker view.

**CDPD-43962: Performance improvement on the Broker Details page**

The `/api/v2/admin/metrics/aggregated/brokers/{brokerId}` endpoint, called every time while opening the broker details page on the UI, was excessively slow when a large number of topics and partitions were present. This is now fixed by fetching partition metrics in bulk.

**CDPD-47836: The FROM OFFSET field of the offset slider in Data Explorer does not update on partition change**

Fixed an issue where the Data Explorer's **FROM OFFSET** field was not updating on partition change.

## Fixed Issues in Apache Tez

Review the list of Tez issues that are resolved in Cloudera Runtime 7.1.9.

**CDPD-48031: Tez - Upgrade jettison to 1.5.3 due to CVE-2022-45685 and CVE-2022-45693**

Upgraded the jettison version to 1.5.3 to fix CVEs

**CDPD-53825: Tez - Upgrade jettison to 1.5.4 due to CVE-2023-1436**

Upgraded jettison to 1.5.4

## Fixed Issues in Apache YARN

Review the list of YARN issues that are resolved in Cloudera Runtime 7.1.9.

**COMPX-14340: YARN-11490 JMX QueueMetrics breaks after mutable config validation in CS**

Fix: JMX metrics broke after 2 or more configuration validation.

**COMPX-13959: Applications submitted to ambiguous queue fail during recovery if "Specified" Placement Rule is used**

Fixed the issue of app killed, if specified placement is used and rm is restarted while the app is still running.

**COMPX-13773: YARN-11461 NPE in determineMissingParents when the queue is invalid**

Fix NPE log warning when submitting to invalid queue.

**COMPX-14120: Backport YARN-11463: Node Labels root directory creation doesn't have a retry logic**

Retry logic is implemented and backported for root directory creation during RM node label store initialization.

**COMPX-10909: Investigate if placement rules are working fine if username contains dot, and default queue is set to that queue**

Username with dot now will work well with CS placement rules

**COMPX-13554: Backport YARN-10178 to 7.1.9 CHEx : Crash in global async scheduler thread**

With this fix the Capacity Scheduler Global Scheduler AsyncThread won't crash when multi async thread concurrently compares queue usage statistics and ResourceCommitterService applies leaf queue change statistics.

**COMPX-12661: YARN-11075 Explicitly declare serialVersionUID in LogMutation class**

The serialVersionUID field is explicitly set for the LogMutation class.

**COMPX-13392: HADOOP-18602 Remove netty3 dependency - CDH-7.1.9**

netty3 is removed

**COMPX-12815: Backport YARN-10178 to 7.1.8 CHEx : Crash in global async scheduler thread**

With this fix the Capacity Scheduler Global Scheduler AsyncThread won't crash when multi async thread concurrently compares queue usage statistics and ResourceCommitterService applies leaf queue change statistics.

**COMPX-12783: Backport YARN-11079 (Make an AbstractParentQueue to store common ParentQueue and ManagedParentQueue functionality)**

Made an AbstractParentQueue to store common ParentQueue and ManagedParentQueue functionality

**COMPX-14124: Backport YARN-10739 GenericEventHandler.printEventQueueDetails causes RM recovery to take too much time**

GenericEventHandler.printEventQueueDetails causes RM recovery to take too much time so added thread pool for async print event details ,to prevent wasting too much time for RM.

**COMPX-14122: Backport YARN-11286: Make AsyncDispatcher#printEventDetailsExecutor thread pool parameter configurable**

Made AsyncDispatcher#printEventDetailsExecutor thread pool parameter configurable

**CDPD-41982: Yarn - Upgrade Guava: Google Core Libraries for Java to v28.2/31.1-jre due to CVEs**

Upgraded Guava Google Core Libraries for Java to v28.2 due to CVEs

**CDPD-57948: [7.1.9 ZDU Simulation] Hive Query is failing when YARN is into rolling restart**

YARN-side fix is implemented and backported to cdpd-master and 7.1.9.x

**COMPX-6054: PlacementPolicy Rules(default rule) is not honoured in case limit 2 is breached for AQC**

This issue is resolved.

**COMPX-5244: Root queue should not be enabled for auto-queue creation**

This issue is resolved.

**COMPX-3181: Application logs does not work for AZURE and AWS cluster**

Support of automatically fetching Delegation Token for YARN Log Aggregation Path (S3 or Azure) in YarnClient.

**OPSAPS-52066: Stacks under Logs Directory for Hadoop daemons are not accessible from Knox Gateway.**

Issue was due to wrong URL being displayed. Both jstacks log viewer and download URLs have been fixed.

**OPSAPS-57067: Yarn Service in Cloudera Manager reports stale configuration yarn.cluster.scaling.recommendation.enable.**

This issue is resolved.

**CDPD-2936: Application logs are not accessible in WebUI2 or Cloudera Manager**

This issue is resolved.

**OPSAPS-50291: Environment variables HADOOP\_HOME, PATH, LANG, and TZ are not getting whitelisted**

"HADOOP\_HOME,PATH,LANG,TZ" are now added by default to the yarn.nodemanager.env-whitelist Yarn configuration option.

**COMPX-3303: Auto queue deletion is not supported in relative and absolute resource allocation mode**

This issue is resolved.

**OPSAPS-68058: [CKP-4] YARN allowed system users are hardcoded**

Allowed system users are now generated dynamically, based on the Kerberos principals, process users and auth-to-local rules.

**OPSAPS-67682: [CKP-3, 4(unequal)] Yarn failed to start the resource manager**

The permissions of the node label directory were eased to allow the process users group members to access it.

**OPSAPS-67860: [BLOCKER] 718CHF9 to 719 | During rolling upgrade Delete the confstore on YARN Zookeeper nodes failed**

The script was fixed to use Kerberos auth instead of relying on digest.

**OPSAPS-68108: Upgrade failures from CDH6 to 7.1.9 because ACL is not the expected for znode after OPSAPS-67993**

Fixed issue with the ACL validator.

**OPSAPS-67993: Upgrade failures from CDH6 to 7.1.9 because ACL is not the expected for znode after OPSAPS-63187**

The bash script was updated to work in a secured environment.

**Apache patch information**

- MAPREDUCE-7237
- MAPREDUCE-7268
- MAPREDUCE-7434
- MAPREDUCE-7433
- MAPREDUCE-7431
- YARN-10930
- YARN-11286
- YARN-10739
- YARN-10178
- HADOOP-18602
- YARN-11190
- YARN-11463
- YARN-11461
- YARN-11513
- YARN-10888
- YARN-11533
- YARN-11490

**Fixed Issues in Zeppelin**

Review the list of Zeppelin issues that are resolved in Cloudera Runtime 7.1.9.

**OPSAPS-68072: Initialize zeppelin notebook command fails on 7.1.9 rhel9.1 cluster with dell isilon**

You can now bring up zeppelin service on 7.1.9 RHEL 9 Isilon cluster.

**CDPD-54867: Backport ZEPPELIN-5176**

[ZEPPELIN-5176] Kerberos ticket renewal fix in JDBC Interpreter.

**CDPD-49444: Backport ZEPPELIN-5624: Arbitrary file deletion vulnerability**

[ZEPPELIN-5624] Arbitrary file deletion vulnerability.

**CDPD-53819: Increase default Zeppelin RPC connection pool size based on ZEPPELIN-5005**

Exposed `zeppelin.interpreter.connection.poolsize` and made it configurable as a safety value in Zeppelin Configuration.

**Apache patch information**

- ZEPPELIN-5176
- ZEPPELIN-5624
- ZEPPELIN-5005

**Fixed Issues in Zookeeper**

Review the list of ZooKeeper issues that are resolved in Cloudera Runtime 7.1.9.

**OPSAPS-67223: ZDU | CDH 717sp2 to 719 AND 718 to 719 | Zookeeper is going in bad health and blocking the upgrade**



Convert the avgRequestLatency metric to double for backward compatibility with Zookeeper versions before 3.8.

### Apache Patch Information

Apache Zookeeper rebased to 3.8.1 and Apache Curator rebased to 5.4.0.

## Known issues in Cloudera Runtime 7.1.9

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.1.9.



**Note:** CDSW does not support RPM-based installation on CDP Private Base. (RPM installation is deprecated and only supported on HDP and CDH 5. For CDH6 and onward, Cloudera recommends you to use CSD-based installations.)

### QAINFRA-18183: CDP Private Cloud Base 7.1.9 does not support KTS on RHEL 9.

None.

**OPSAPS-69539: CDP Runtime 7.1.9 from the base release through to CHF3 does not support Oracle JDK 8u401 or OpenJDK 1.8.0\_402 (8u402). Some services will fail to start. This can be a problem on RHEL 9.x as version 8u402 is the default OpenJDK 8 installed by the OS.**

Workaround is to install an earlier version of JDK 8. For example Oracle jdk-8u291 / 1.8.0\_291, or OpenJDK 8u292 / 1.8.0\_292.

## Known Issues in Apache Atlas

Learn about the known issues in Atlas, the impact or changes to the functionality, and the workaround.

**CDPD-6565: Issue with ddlQueries in Atlas for Hive/Impala tables created in default/non-default db using DAS/HUE respectively**

ddlQueries are not created for Impala origin table. ddlQueries created for Impala/Hive CTAS tables have difference in names as compared to same operation for Hive table. The Impala table has the query in the ddlQueries entity name.

None

**CDPD-55301: ddlQueries and ALTERNATIVE\_\* lineage missing for Spark tables created through spark3-shell**

The ddlQueries and ALTERNATIVE\_\* lineage missing for Spark tables created using spark3-shell.

None

**CDPD-56085: 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**

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\*.

None

**CDPD-58554: Discard audits of specific classification, label, and business metadata**

Support to control audits for specific classification, label, and business metadata is not present in the Custom Audit Filters feature.

None

**CDPD-59586: OperationType does not support Contains operation.**

Support for sub-string and wild character '\*' is not present for attributeName=operationType in Custom Audit Filters feature, hence attributeValue can not be checked with "contains" operator.

None

**OPSAPS-67783: During rolling upgrade one among two Atlas server failed to start but Cloudera Manager considered as success**

Cloudera Manager marks the Execute command Start on service Atlas-1 as a success even when Atlas service had failed to start successfully. In such cases, Atlas logs give the exact reason for Atlas start-up failure.

None

**CDPD-19358: "IsIndexable" and "isOptional" value of a typedef's attribute is modified post migration.**

In HDP-265, falcon\_feed\_creation has "stored-in" attribute which has IsIndexable value set to True.

Post migration, "stored-in" is moved to relationshipAttributeDefs and has IsIndexable value set to False.

Similarly, in HDP-265 the table attribute of hive\_storage\_desc isOptional is set to true. Post migration, isOptional set to False.

None.

**CDPD-11941: Table creation events missed when multiple tables are created in the same Hive command**

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**

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**

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**

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**

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**

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) does not 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 the metadata collection from HMS. For example, if an ALTER TABLE operation in Spark that is changing a table name and is reporting 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 "@" does not 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, and use 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 does 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 logout 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 CDP.

None.

**CDPD-1664: Guest users are redirected incorrectly**

Authenticated users logging in to Atlas are redirected to the CDP 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: The 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, to add metadata tags to relationships that you wanted to make sure were unique in the system, you can design the relationship attribute with the property `IsUnique` equal `true`. However, in this release, the `IsUnique` attribute is not enforced.

None.

**CDPD-39427: [HDFS Lineage]: When the input is a directory in case of put/copyfromLocal/cp/mv commands, lineage is not created even though the script succeeds.**

When Source is a directory and target is a directory which is already present in Atlas, the command succeeds and inserts the data in the desired location, but lineage is not created.

**DOCS-13759: Tag Propagation stops after a certain depth while the lineage is being extended**

When a tag is added to an entity at timestamp T1, the entities along the lineage to which the tag must be propagated is calculated at T1. Before tag propagation completes, if the lineage is extended, tag does not propagate to the entities in the extended lineage.

**DOCS-13760: System Attributes search, \_\_classificationNames: Search with parent tag does not return entities associated to its children tags**

System attribute search with `__classificationNames = parent_tag` returns entities associated to `parent_tag` only and not entities associated to its children tag.

Workaround: Instead of using system attribute, employ the basic search attribute "classification" which lists entities associated with inherited classifications.

**CDPD-41142: When a Kafka console consumer group is run, more than one update audits are seen**

After running the console consumer with a consumer group, verify the consumer group entity created, along with the metrics and notifications for the consumer group and topic. The expected result can be: one `ENTITY_CREATE` audit and one `ENTITY_UPDATE` audit. But more than one `ENTITY_UPDATE` audits are seen.

**CDPD-40165: Two audits are created for SPARK CTAS table**

When following Spark queries are fired:

```
spark.sql("create table table1(id int)")
spark.sql("create table table2 as select * from table1")
```

HMS sends "ENTITY\_CREATE" and "ENTITY\_FULL\_UPDATE\_V2".

The extra `ENTITY_FULL_UPDATE_V2` message received from HMS is sent as a part of `ALTE RTABLE_ADDCOLS` event from the HMS Hook side. This behaviour is observed only when the queries are run from Spark SQL and not when run the same queries are run from Beeline.

**CDPD-39197: Debug metrics returns empty data**

When debug metrics is enabled and some operations are performed, the response is empty

**CDPD-36495: Updating legacyAttribute from False to True resets the initially created relationshipAttributes values**

Creating types, entities, and to start, you must set the relationship with `is_legacy_attribute` value as **False**.

Later, update the value `relationshipDef is_legacy_attribute` to **True**.

For the entities that were created before updating the `is_legacy_attribute` to **True**, the `relationshipAttributes` value is reset.

**CDPD-35818: Basic search with tag filter provides approximateCount as -1 when there is no match and is 0 otherwise**

When the following search operations are performed:

- Faceted search with both tag filter and entity filter

The observed approximateCount is -1.

Here both entity filter and tag filter are present and when there is no match the response received is -1.

- Faceted search with only entity filter

Performing a basic search provides an approximate value of 0 when there is no match.

- Faceted search with only tag filter

Whenever there is tag filter in the query and there is no entity match, the approximateCount is -1 and if the tag filter is not available, the response approximateCount is 0

**CDPD-13466: Bulk create/update entity POST API does not create / update authorised entities**

The bulk API fails with 403 error if some belong to entities on which the user is unauthorized and other GUIDs belong to entities on which user is authorized.

**CDPD-22744: Bulk entity DELETE API does not delete authorised entities**

Bulk entity DELETE API does not delete authorised entities when the list of authorised and unauthorised entities list is passed.

**CDPD-29409: Hive import: Suggestion suggests entity which is deleted.**

Suggestions suggests tables of a database, which is a deleted entity.

**CDPD-25152: Tag propagation through deferred actions consumes additional time as compared to default flow**

The additional time might be due to the small overhead added to create / update task vertex and which is run in the background. This also depends on number of tasks queued to be executed.

**CDPD-42954: Zeppelin notebook fails after enabling Atlas-HDFS hook**

The Zeppelin notebooks are failing with errors after enabling Atlas-HDFS hook in the CDP cluster.

When the below properties are set for atlas-client.properties in Cloudera Manager:

- atlas.jaas.KafkaClient.option.keyTab
- atlas.jaas.KafkaClient.option.principal

Along with adding the properties in /etc/atlas/conf/atlas-application.properties, Cloudera Manager also adds these properties to atlas-application.properties for other services (like Spark).

Adding these properties interferes with the normal flow of the services (like Spark)

To enable HDFS lineage feature, instead of setting these properties through Cloudera Manager, users can manually add the properties directly in /etc/atlas/conf/atlas-application.properties

**CDPD-40346: The ddlQueries and ALTERNATIVE\_ADDCOLS lineage missing for Impala tables.**

The ALTERNATIVE\_ADDCOLS lineage has some issue when an Impala table is altered and the corresponding lineage is not created.

**CDPD-62464: Java process called by navatlas.sh tool fails on JDK-8 version**

While running nav2atlas.sh script on OracleJDK 8 an error message is thrown and returns code 0 on an unsuccessful run.

You must install JDK-11 version on the host. Make sure not to put into the default path and JAVA\_HOME. In a shell, set the JAVA\_HOME to this location and run the nav2atlas.sh script.

**CDPD-62834: Status of the deleted table is seen as ACTIVE in Atlas after the completion of navigator2atlas migration process**

The status of the deleted table displays as ACTIVE.

None

**CDPD-62837: During the navigator2atlas process, the hive\_storagedesc is incomplete in Atlas**

For the hive\_storagedesc entity, some of the attributes are not getting populated.

None

**CDPD-74467: Ranger TagSync does not support ofs path parsing for Ozone-Atlas**

When trying to create Hive table on top of Ozone ofs path with the command ofs://ozone1/volume1/bucket1/table\_oz\_demo1, Ranger TagSync fails with an error.

Upgrade to Cloudera Private Cloud Base 7.1.9 SP1 CHF3.

**CDPD-79160: NPE while deleting BusinessMetadata**

If business metadata is created without adding any applicable types, a NullPointerException is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

**CDPD-76035: Resource lookup for Atlas service is failing**

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

**CDPD-76789: Creating tag with name description throws java.lang.ClassCastException**

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-81499: Basic search is lagging as some Datatypes not getting indexed in Solr**

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-76789: Creating tag with name description throws java.lang.ClassCastException**

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails**

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

**CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeUpException: null" is seen**

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

**CDPD-87670: [Atlas UI] Apache Atlas Glossary becomes unresponsive when the page size is set to 50**

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-28741: Atlas node start failure : PermanentLockingException: Expected value mismatch for KeyColumn**

In an Atlas High Availability (HA) enabled environment, starting both Atlas hosts at the same time can result in both trying to access or modify the same graph database resulting in a race condition. One instance successfully modifies the patch vertex, while the other instance triggers a Locking Exception (Concurrent modification exception).

Start Atlas instances in sequence and not in parallel to avoid the potential race condition. If the exception is already present, restart the second Atlas instance after triggering the PermanentLockingException.

#### **CDPD-79099: Unable to import hive tables using import script in Public Cloud**

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.

### **Known Issues in Apache Avro**

Learn about the known issues in Avro, the impact or changes to the functionality, and the workaround.

**CDPD-23451: Avro library depends on the already EOL jackson-mapper-asl 1.9.13-cloudera.1 that also contains a couple of CVEs. The jackson library is part of the Avro API so cannot be changed without a complete rebase.**

None.

### **Known issues in Cloud Connectors**

Learn about the known issues in Cloud Connectors, the impact or changes to the functionality, and the workaround.

**OPSAPS-72509, CDPD-32440: Hive metadata transfer to GCS fails with ClassNotFoundException**

Hive replication policies from an on-premises cluster to cloud fails during the “Transfer Metadata Files” step if the following conditions are true:

- the target is a GCS Data Lake
- the source Cloudera Manager version is 7.11.3 CHF7, 7.11.3 CHF8, 7.11.3 CHF9, 7.11.3 CHF9.1, 7.11.3 CHF10, or 7.11.3 CHF11

This is because the fs.gs.delegation.token.binding property is already defined in the configuration and cannot be unset to disable the delegation tokens in the cloud connector service.

None

### **Known issues in Cruise Control**

Learn about the known issues in Cruise Control, the impact or changes to the functionality, and the workaround.

**Rebalancing with Cruise Control does not work due to the metric reporter failing to report the CPU usage metric**

On the Kafka broker, the Cruise control metric reporter plugin may fail to report the CPU usage metric.

If the CPU usage metric is not reported, the numValidWindows in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.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 the [Moving a Host Between Clusters](#).



**Note:** Cluster nodes affected by this issue are not displayed as unhealthy.

### OPSAPS-69978: Cruise Control capacity.py script fails on Python 3

Cruise Control might fail to start on Python 3 when capacity information is queried during the startup process. This is caused by a breaking change between Python 2 and 3.

None

## Known Issues in Apache Hadoop

There are no known issues for Hadoop in Cloudera Runtime 7.1.9.

## Known Issues in Apache HBase

This topic describes known issues and workarounds for using HBase in this release of Cloudera Runtime.

### HBase throws a NullPointerException on MemstoreFlusher when flush is triggered by too many WALs after a WAL rolling

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.

### In the affected CDP releases (7.1.7.1056+, 7.1.7 SP2+, 7.1.8 CHF1+, 7.2.16.0+, 7.2.15.2+, 7.2.14.3+), Hbase Regionserver performance is degraded by aggressive GC pause and memory contention from the ColumnFamilyCache.

Disable ColumnFamilyCache Cloudera Manager -> HBase -> HBase Service Advanced Configuration Snippet (Safety Valve) for `ranger-hbase-security.xml` `xasecure.hbase.columnfamilies.cache.maxsize=0`

### CDPD-44402: If you are running CDP 7.1.8 and >7.1.7.1056 SP1 releases, there is an increased heap usage and GC pressure in the HBase regionserver due to the Ranger ColumnFamily Cache feature. This problem is particularly evident when HBase tables that contain numerous columns, leading to high heap utilization and potentially causing the Regionserver to fail.

Disable the Ranger ColumnFamily Cache in Cloudera Manager by setting `xasecure.hbase.columnfamilies.cache.maxsize` to 0 in `ranger-hbase-security.xml`. This cache can be safely turned off, as it reverts to the behavior seen in version 7.1.7 without any impact.

### IntegrationTestReplication fails if replication does not finish before the verify phase begins



During `IntegrationTestReplication`, if the verify phase starts before the replication phase finishes, the test will fail because the target cluster does not contain all of the data. If the HBase services in the target cluster does not have enough memory, long garbage-collection pauses might occur.

Use the `-t` flag to set the timeout value before starting verification.

### HDFS encryption with HBase

Cloudera has tested the performance impact of using HDFS encryption with HBase. The overall overhead of HDFS encryption on HBase performance is in the range of 3 to 4% for both read and update workloads. Scan performance has not been thoroughly tested.

N/A

### AccessController postOperation problems in asynchronous operations

When security and Access Control are enabled, the following problems occur:

- If a Delete Table fails for a reason other than missing permissions, the access rights are removed but the table may still exist and may be used again.
- If `hbaseAdmin.modifyTable()` is used to delete column families, the rights are not removed from the Access Control List (ACL) table. The `portOperation` is implemented only for `postDeleteColumn()`.
- If Create Table fails, full rights for that table persist for the user who attempted to create it. If another user later succeeds in creating the table, the user who made the failed attempt still has the full rights.

N/A

Apache Issue: [HBASE-6992](#)

### Snappy compression with /tmp directory mounted with noexec option

Using the HBase client applications such as `hbase hfile` on the cluster with Snappy compression could result in `UnsatisfiedLinkError`.

Add `-Dorg.xerial.snappy.tmpdir=/var/hbase/snappy-tmpdir` to Client Java Configuration Options in Cloudera Manager that points to a directory where `exec` option is allowed.

### HBase shutdown can lead to inconsistencies in META

Cloudera Manager uses an incorrect shutdown command. This prevents graceful shutdown of the HBase service and forces Cloudera Manager to kill the processes instead. It can lead to inconsistencies in Meta.

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.

### CDPD-77399: HBase fails to register the servlet metrics and throws `ClassNotFoundException: org.apache.hadoop.metrics.MetricsServlet`

The `MetricsServlet` class is a Hadoop 2-based metric servlet unavailable in Hadoop 3 deployments.

Workaround: Ignore this WARN log message during HBase Master and RegionServer startup.

Apache Issue: [HBASE-28315](#)

## Known Issues in HDFS

Learn about the known issues in HDFS, the impact or changes to the functionality, and the workaround.

### 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 the client has attempted to use all the datanodes in the cluster, and failed to write to each of them as they were restarted. This will only happen on small clusters of less than 10 datanodes, as larger clusters have more spare node to allow the write to continue.

None.

### CDPD-60873: java.io.IOException: Got error, 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      | <code>dfs.permissions.ContentSummary.subAccess</code> | Not Set   | True      | Performance optimization for NN content summary API   |
| HDFS      | <code>dfs.datanode.handler.count</code>               | 3         | 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      | <code>dfs.namenode.access.time.precision</code> | None      | 0         | Optimal value for NN performance on large clusters    |
| HDFS      | <code>dfs.datanode.handler.count</code>         | 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 bootstrap 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. Try to add the new NameNode again



**Note:** Note that entering safe mode will disable 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 try adding 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-60832: When decommission of DN runs for a longer time and when decommission monitor's kerberos ticket expires, it is not auto-renewed. Decommission of DN is not completed in Cloudera Manager as decommission monitor fails to fetch the state of DN after kerberos ticket expiry.**

Decommission state of DN can be fetched using CLI command `hdfs dfsadmin -report`.

**OPSAPS-55788: WebHDFS is always enabled. The Enable WebHDFS checkbox does not take effect.**

None.

**OPSAPS-63299: Disable HA command for a nameservice does not work if the nameservice has more than 2 NNs defined.**

None

**OPSAPS-63301: Deleting nameservice command does not delete all the NNs belonging to the nameservice, if there are more than two NNs that are assigned to the nameservice.**

None

### Unsupported Features

The following HDFS features are currently not supported in Cloudera Data Platform:

- ACLs for the NFS gateway ([HADOOP-11004](#))
- Aliyun Cloud Connector ([HADOOP-12756](#))
- Allow HDFS block replicas to be provided by an external storage system ([HDFS-9806](#))
- Consistent standby Serving reads ([HDFS-12943](#))
- Cost-Based RPC FairCallQueue ([HDFS-14403](#))
- HDFS Router Based Federation ([HDFS-10467](#))
- NameNode Federation ([HDFS-1052](#))
- NameNode Port-based Selective Encryption ([HDFS-13541](#))
- Non-Volatile Storage Class Memory (SCM) in HDFS Cache Directives ([HDFS-13762](#))
- OpenStack Swift ([HADOOP-8545](#))
- SFTP FileSystem ([HADOOP-5732](#))
- Storage policy satisfier ([HDFS-10285](#))

## Known Issues in Apache Hive

Learn about the known issues in Hive, the impact or changes to the functionality, and the workaround.

**CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress**

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL

- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

#### CDPD-67002: Merge/Update jobs fail with vertex errors after upgrading to 7.1.9

In CDP 7.1.9, improvements for update and delete operations were introduced to write data directly to table locations instead of using a staging directory. However, after upgrading to CDP 7.1.9, merge, update, and delete queries may occasionally fail with the error:

```
java.lang.RuntimeException: Hive Runtime Error while closing operators: Index 17 out of bounds for length 17
 at org.apache.hadoop.hive ql.exec.tez.TezProcessor.initializeAndRunProcessor(TezProcessor.java:298)
 at org.apache.hadoop.hive ql.exec.tez.TezProcessor.run(TezProcessor.java:252)
 at org.apache.tez.runtime.LogicalIOProcessorRuntimeTask.run(LogicalIOProcessorRuntimeTask.java:374)

```

```
Caused by: java.lang.RuntimeException: Hive Runtime Error while closing operators: Index 17 out of bounds for length 17
 at org.apache.hadoop.hive ql.exec.tez.ReduceRecordProcessor.close(ReduceRecordProcessor.java:407)

```

```
Caused by: java.lang.ArrayIndexOutOfBoundsException: Index 17 out of bounds for length 17
 at org.apache.hadoop.hive ql.exec.FileSinkOperator$FSPaths.closeWriters(FileSinkOperator.java:258)
 at org.apache.hadoop.hive ql.exec.FileSinkOperator.closeOp(FileSinkOperator.java:1483)
 at org.apache.hadoop.hive ql.exec.Operator.close(Operator.java:731)
 at org.apache.hadoop.hive ql.exec.Operator.close(Operator.java:755)
 at org.apache.hadoop.hive ql.exec.tez.ReduceRecordProcessor.close(ReduceRecordProcessor.java:383)
```

To resolve this issue, disable the direct insert, update, and delete features by setting the following parameters at session level:

```
set hive.acid.direct.insert.enabled=false;
set hive.acid.direct.update.and.delete.enabled=false;
```



#### Note:

- When running scheduled queries through an application or script, ensure the relevant configurations are added to the application's settings or script as needed.
- Ensure these parameters are whitelisted to allow runtime configuration.

#### CDPD-60770: Beeline Authentication Issue with Special Characters in Passwords

When LDAP is enabled, users cannot authenticate with Beeline if the password contains a special character. For example, the following string fails:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>;user=user@XXX;password='R3G#xpXyoy1M0Jb1'
```

Use the `-p` parameter to execute the Beeline command:

```
beeline -u jdbc:hive2://<host>:<port>/<dbName>; -n user@XXX -p 'R3G#xpXyoy1M0Jb1'
```

#### CDPD-43769: Inconsistent behavior with CHAR comparisons and string literals

In Hive, 'char' data to 'string' data comparison is considered as equal, but 'char' data to 'varchar' data comparison is considered as not equal if the data has trailing spaces. For example,

```
create table test(colchar char(10), colvarchar varchar(20), colstring string);
insert into test values ('a ', 'a ', 'a ');

select count(*) from test where colchar = colstring;
Output returns 1 because char and string data is considered as equal

select count(*) from test where colchar = colvarchar;
Output returns 0 because char and varchar data is considered as not equal
```

It is recommended that you use 'varchar or string' data type instead of 'char' data type for any character or string representation.

#### CDPD-44060: Issue rebuilding Hive Materialized Views

The `ALTER MATERIALIZED VIEW <view_name> REBUILD;` command to rebuild Materialized View fails with the following error:

```
Error: Error while compiling statement: FAILED: SemanticException
org.apache.hadoop.hive.ql.parse.SemanticException: Another process
is rebuilding the materialized view activecore.top_apps_custom
(state=42000,code=40000)
```

This can occur if the Materialized View Rebuild command is abruptly killed or if more than one Materialized View Rebuild command is issued for the same Materialized View. As a result, LOCK entry in the backend metastore DB table 'materialization\_rebuild\_locks' is not deleted and when the command is issued again, it will fail for that Materialized View.

To resolve this issue, perform the following steps to manually delete the stale locks:

1. From the backend DB, run the following command and check if there are any entries for the view to confirm if stale locks are present:

```
select mrl_db_name, mrl_tbl_name, mrl_last_heartbeat from materialization_rebuild_locks where lower(mrl_tbl_name)='<view_name>';
```

2. Create a backup of the table if there are multiple entries:

```
create table bkp_top_apps_custom as
select mrl_db_name, mrl_tbl_name, mrl_last_heartbeat from materialization_rebuild_locks where lower(mrl_tbl_name)='<view_name>';
```

3. Delete the stale lock entries from the table:

```
delete from materialization_rebuild_locks where lower(mrl_tbl_name)='<view_name>';
```

For more information, see the related [Cloudera Community article](#).

#### **CDPD-42726: Wrong results for agg queries from stats after running multi-insert query on managed table**

Multiple inserts at the same time into the same partition results in invalid stats. For example:

```
create table source(p int, key int, value string);
insert into source(p, key, value) values (101,42,'string42');

create table stats_part(key int,value string) partitioned by (p
int);
from source
insert into stats_part select key, value, p
insert into stats_part select key, value, p;

select count(*) from stats_part;
-- In this scenario, StatsOptimizer helps serving this query be
cause the result should be rowNum of the partition p=101.
-- The result is 1, however, the result should be 2.
```

Rewrite the multi-insert statement by merging the branches inserting into the same partition to one branch.

#### **CDPD-57574: Query execution fails due to NullPointerException in DagUtils.setupQuickStart**

Hive queries may fail at runtime with a NullPointerException while executing the Tez graph at DagUtils.setupQuickStart method. A part of the stack trace from the error is shown below:

```
ERROR : Failed to execute tez graph.
java.lang.NullPointerException: null at org.apache.hadoop.hive.ql.exec.tez.DagUtils.setupQuickStart(DagUtils.java:1724)
```

The error usually occurs from an invalid plan that has cycles created by semi-join or map-join edges.

Disable some of the optimizations in order to remove the cycles from the plan. It is sufficient to set one of the following properties:

- set hive.tez.dynamic.partition.pruning=false;
- set hive.tez.dynamic.semijoin.reduction=false;
- set hive.auto.convert.join=false;

To avoid any undesired performance regressions, disable the optimization(s) only for specific queries and not globally (through Cloudera Manager) for the whole workload or cluster.

#### **CDPD-43107: SemanticException for INSERT INTO statement when Hive Cost-based Optimizer (CBO) is disabled**

If you are running the INSERT INTO query and hive.cbo.enable is set to "false", the query fails with a SemanticException. For example,

```
set hive.cbo.enable=false;

CREATE TABLE mytable (
 id INT,
 str STRING
);
```

```
INSERT INTO mytable (id, str) VALUES (1, 'a');
```

```
Output:
org.apache.hadoop.hive.ql.parse.SemanticException: 0:0 Expected
2 columns for insclause-0/default@mytable; select produces 1 col
umns. Error encountered near token ''a''
```

This issue occurs only when columns are specified for the target table in the query. INSERT INTO mytable values (1, 'a'); does not result in an exception.

Enable CBO and run the query.

#### CDPD-43957: HiveServer shuts down during replication due to high resource usage

During Hive replication (both bootstrap and incremental), you may notice that the HiveServer (HS2) shuts down periodically with the following error:

```
java.sql.SQLException: org.apache.hive.jdbc.ZooKeeperHiveClientE
xception: Unable to read HiveServer2 configs from ZooKeeper
 at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.ja
va:265)
 at org.apache.hive.jdbc.HiveDriver.connect(HiveDriver.java:107)
 at java.sql.DriverManager.getConnection(DriverManager.java:664)
 at java.sql.DriverManager.getConnection(DriverManager.java:247)
 at com.cloudera.enterprise.hive3qt.Hive3QueryTool$HiveOperation.
execute(Hive3QueryTool.java:682)
 at com.cloudera.enterprise.hive3qt.Hive3QueryTool.main(Hive3Qu
eryTool.java:935)
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException:
Unable to read HiveServer2 configs from ZooKeeper
 at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureCo
nnParams(ZooKeeperHiveClientHelper.java:177)
 at org.apache.hive.jdbc.Utills.configureConnParamsFromZooKeeper
(Utills.java:580)
 at org.apache.hive.jdbc.Utills.parseURL(Utills.java:391)
 at org.apache.hive.jdbc.HiveConnection.<init>(HiveConnection.j
ava:263)
 ... 5 more
Caused by: org.apache.hive.jdbc.ZooKeeperHiveClientException: T
ried all existing HiveServer2 uris from ZooKeeper.
 at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.getServerHosts
(ZooKeeperHiveClientHelper.java:132)
 at org.apache.hive.jdbc.ZooKeeperHiveClientHelper.configureCon
nParams(ZooKeeperHiveClientHelper.java:172)
 ... 8 mor
```

This issue occurs when you replicate at scale with an unbalanced cluster setup that has all the roles running on the Cloudera Manager host. As a result, Cloudera Manager ends up in a bottleneck situation because HS2 crashes, preventing further replication. The logs (/var/log/messages) indicate that there were a large number of Java processes running on the host, which exhausted the host's memory and triggered an Out Of Memory Killer process to stop HS2.

You can restart HS2 and reinitiate replication. However, the replication process may take a longer time to complete.

- Follow these recommendations when setting up your source and target clusters:
  - Ensure that no HS2, Hive metastore (HMS), DataNode, or NameNode roles are running on the Cloudera Manager host.
  - Ensure that you have multiple instances of HS2 and HMS roles on different nodes.
  - It is recommended that you have the HS2 and HMS roles on different nodes than the DataNodes or NameNodes.

These practices increase the possibility of Hive replication at scale completing successfully.

- Limit the scale of Hive replication by temporarily disabling replication policies.

#### **CDPD-40730: Parquet change can cause incompatibility**

Parquet files written by the parquet-mr library in CDP 7.1.9, where the schema contains a timestamp with no UTC conversion will not be compatible with older versions of Parquet readers. The effect is that the older versions will still consider these timestamps as they would require UTC conversions and will thus end up with a wrong result. You can encounter this problem only when you write Parquet-based tables using Hive, and tables have the non-default configuration `hive.parquet.write.int64.timestamp=true`.

None.

#### **CDPD-41274: HWC + Oozie issue: Could not open client transport with JDBC Uri**

Currently only Spark cluster mode is supported in the Oozie Spark Action with Hive Warehouse Connector (HWC).

Use Spark action in cluster mode.

```
<spark xmlns="uri:oozie:spark-action:1.0">
 ...
 <mode>cluster</mode>
 ...
</spark>
```

#### **CDPD-26556 After an upgrade, querying a CTAS table under certain conditions might throw an exception**

If you upgrade your Hive cluster from CDH 6 to CDP 7, create a CTAS table in the CDP cluster from a table you upgraded from CDH, you might see the following exception when you query the new table:

```
class org.apache.hadoop.io.IntWritable cannot be cast to class o
rg.apache.hadoop.hive.serde2.objectinspector.StandardUnionObject
Inspector$StandardUnion
```

This issue involves CDH-based tables having columns of complex types ARRAY, MAP, and STRUCT.

#### **CDPD-23506: OutOfMemoryError in LLAP**

Long running spark-shell applications can leave sessions in interactive Hiveserver2 until the Spark application finishes (user exists from spark-shell), causing memory pressure in case of a high number of queries in the same shell (1000+).

You must close spark-shell so that sessions are closed. Add the owner of the database or the tables as a user with read or read/write access to the tables directly.

#### **CDPD-23041: DROP TABLE on a table having an index does not work**

If you migrate a Hive table to CDP having an index, DROP TABLE does not drop the table. Hive no longer supports indexes ([HIVE-18448](#)). A foreign key constraint on the indexed table prevents dropping the table. Attempting to drop such a table results in the following error:

```
java.sql.BatchUpdateException: Cannot delete or update a parent
row: a foreign key constraint fails ("hive"."IDXS", CONSTRAINT "
```



```
IDXS_FK1" FOREIGN KEY ("ORIG_TBL_ID") REFERENCES "TBLS ("TBL_ID"
))
```

There are two workarounds:

- Drop the foreign key "IDXS\_FK1" on the "IDXS" table within the metastore. You can also manually drop indexes, but do not cascade any drops because the IDXS table includes references to "TBLS".
- Launch an older version of Hive, such as Hive 2.3 that includes IDXS in the DDL, and then drop the indexes as described in [Language Manual Indexing](#).

Apache Issue: [Hive-24815](#)

#### **CDPD-17766: Queries fail when using spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive.**

Do not specify spark.sql.hive.hiveserver2.jdbc.url.principal in the JDBC URL to invoke Hive remotely.

Workaround: specify principal=hive.server2.authentication.kerberos.principal as shown in the following syntax:

```
jdbc:hive://<host>:<port>/<dbName>;principal=hive.server2.authen
tication.kerberos.principal;<otherSessionConfs>?<hiveConfs>#<hive
eVars>
```

#### **CDPD-13636: Hive job fails with OutOfMemory exception in the Azure DE cluster**

Set the parameter hive.optimize.sort.dynamic.partition.threshold=0. Add this parameter in Cloudera Manager (Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml)

#### **CDPD-10848: HiveServer Web UI displays incorrect data**

If you enabled auto-TLS for TLS encryption, the HiveServer2 Web UI does not display the correct data in the following tables: Active Sessions, Open Queries, Last Max n Closed Queries

#### **OPSAPS-72159: Data corruption issue in PostgreSQL databases after OS upgrade**

Upgrading to CentOS 8 or Red Hat Enterprise Linux (RHEL) 8 results in PostgreSQL database corruption due to changes in the GNU C Library (glibc). The corruption affects B-Tree indexes on text columns, causing Hive metadata issues, duplicate tables, and failed drop table operations.

For details and resolution steps, see the Knowledge Base article [Data corruption issue in PostgreSQL databases after OS upgrade to CentOS 8 or RHEL 8](#).

### **Technical Service Bulletins**

#### **TSB 2021-501: JOIN queries return wrong result for join keys with large size in Hive**

JOIN queries return wrong results when performing joins on large size keys (larger than 255 bytes). This happens when the fast hash table join algorithm is enabled, which is enabled by default.

#### **Knowledge article**

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2021-501: JOIN queries return wrong result for join keys with large size in Hive](#)

#### **TSB 2023-702: Potential wrong result for queries with date partition filter for clusters in GMT+ timezone**

In Cloudera Data Platform (CDP) Private Cloud Base 7.1.7 Service Pack (SP) 2 Cumulative Hotfix (CHF) 11, a fix was introduced in Hive Metastore (HMS) to address a parsing issue with date strings. This fix caused a regression in Hive clusters where the HMS time zone is set ahead of GMT for the following combination of tables and queries: a table that is partitioned on a DATE column and a SELECT query on that table containing a WHERE clause filter on the same DATE column. For such queries, during the partition pruning phase, the date string would be converted to a date without timezone and compared with the partition value retrieved by HMS. This causes wrong results (0 rows) because the date values do not match.

The regression was identified in CDP Private Cloud Base 7.1.7 SP2 CHF14, but it exists in CHF11 through CHF16 as well as on certain versions of 7.1.8 and 7.1.9.

This issue does not affect clusters where the time zones are behind GMT. For example, if the time zone of the cluster is set USA/Los Angeles, which is 8 hours behind GMT, a date '2023-10-02' will remain as '2023-10-02' after converting to GMT (adding 8 hours). On the other hand, using Asia/Hong Kong time as an example, which is 8 hours ahead of GMT, the same date would become '2023-10-01' after converting to GMT (subtracting 8 hours), which leads to the wrong results.

#### Upstream JIRA

[HIVE-27760](#)

#### Knowledge article

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2023-702: Potential wrong result for queries with date partition filter for clusters in GMT+ timezone](#)

## Known Issues in Hue

Learn about the known issues in Hue, the impact or changes to the functionality, and the workaround.

### **OPSAPS-69659: Hue service fails on restart with "Unable to find psycopg2 2.5.4" error**

Hue service fails to restart and you see the following error: Unable to find psycopg2 2.5.4. This could be because you have installed Python in a non-default location and Hue is unable to locate the psycopg2 PostgreSQL database adapter.

You must specify the path where you have installed Python in the PYTHONPATH property in the Hue Advanced Configuration Snippet using Cloudera Manager.

1. Log in to Cloudera Manager as an Administrator.
2. Go to Clusters Hue Configurations and add the following key and value in the Hue Service Environment Advanced Configuration Snippet (Safety Valve) field:

Key: PYTHONPATH

Value: `[***PYTHON-PATH***]`

Replace `[***PYTHON-PATH***]` with the actual location where you have installed Python. For example, `/opt/cloudera/parcels/CDH/lib/hue/build/env/lib/python3.8/site-packages`

3. Click Save Changes.
4. Restart the Hue service.

### **CDPD-58978: Batch query execution using Hue fails with Kerberos error**

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark queries in a batch mode launches a shell script Oozie job from Hue and this is not supported on a secure cluster.

There is no workaround. You can submit the queries individually.

### **CDPD-59677: Unable to view Phoenix tables on the left assist in Hue**

On clusters secured with Knox, you may not be able to see Phoenix tables on the left assist that are present under the default database (that is, an empty("") database).

None.

### **CDPD-58142: A query is not pre-populated in the Hue editor after clicking on the "Re Execute" button**

When you click Re Execute to rerun a query from the Job Browser Queries Query Details page, the query does not get populated on the Hue editor, as expected.

None.

### **CDPD-54376: Clicking the home button on the File Browser page redirects to HDFS user directory**

When you are previewing a file on any supported filesystem, such as S3 or ABFS, and you click on the Home button, you are redirected to the HDFS user home directory instead of the user home directory on the said filesystem.

None.

**CDPD-41306: pip3.8 freeze command does not work and results into an error**

You see the `/usr/lib/hue/build/env/bin/python: No such file or directory` error when you run the following command:

```
build/env/bin/pip3.8 freeze
```

Run the freeze command as follows by specifying the paths of python3.8 and pip3.8:

```
/opt/cloudera/parcels/CDH/lib/hue/build/env/bin/python3.8 /opt/cloudera/parcels/CDH/lib/hue/build/env/bin/pip3.8 freeze
```

**CDPD-43293: Unable to import Impala table using Importer**

Creating Impala tables using the Hue Importer may fail.

If you have both Hive and Impala services installed on your cluster, then you can import the table using by selecting the Hive dialect from **Tables Sources**.

If only Impala service is installed on your cluster, then go to **Cloudera Manager Clusters Hue Configurations** and add the following line in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field:

```
[beeswax]
max_number_of_sessions=1
```

**CDPD-41136: Importing files from the local workstation is disabled by default**

Cloudera has disabled the functionality to import files from your local workstation into Hue because it may cause errors. You may not see the Local File option in the Type drop-down menu on the **Importer** page by default.

You can enable the functionality to import files from your local workstation by specifying the following parameter in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field in Hue configurations in Cloudera Manager:

```
[indexer]
enable_direct_upload=true
```

**DWX-8602: Unable to import a large CSV file from the local workstation**

You may see an error message while importing a CSV file into Hue from your workstation, stating that you cannot import files of size more than 200 KB.

Upload the file to S3 or ABFS and then import them into Hue using the Importer.

**CDPD-39330: Unable to use the pip command in CDP**

You may not be able to use the pip command in CDP 7.1.7 or higher and may see the following error when using pip in a command: "ImportError: cannot import name chardet".

Follow the steps listed on [Unable to use pip command in CDP](#).

**CDPD-24294: Hue uses the unsafe-inline directive in its Content Security Policy (CSP) header**

Hue 4 web interface uses the unsafe-inline directive in its CSP header. As a result, the application server does not set the CSP header in its HTTP responses, and therefore does not benefit from the additional protection against potential cross-site scripting issues and other modern application vulnerabilities which a properly configured CSP may provide. This could lead to application vulnerability.

Cloudera recommends deploying additional security measures such as a firewall within the Hue server to control allowed connections, and SSO-based authentications mechanisms such as LDAP or SAML.

#### **OPSAPS-61244: Cloudera Manager displays stale Hue configuration after upgrading to CDP 7.1.x from CDH 6.**

After upgrading from CDH 6 to CDP 7.1.x, you may see stale configurations in Cloudera manager for the Hue service.

Manually restart the Hue service from Cloudera Manager.

#### **ENGESC-9091: Setting idle session timeout for Hue does not work when the cluster is secured using Knox SSO**

If Hue is configured with `desktop.auth.backend.KnoxSpnegoDjangoBackend` as the Authentication Backend, then the automatic idle session logout that is set by configuring the `idle_session_timeout` property does not take effect. You may also see 404 error while accessing Hue from the Knox UI when the `idle_session_timeout` property is not set to -1.

None

#### **DOCS-10377: Hue UI is blank upon login after upgrading to CDP 7.1.7 from CDH 6**

If your cluster was secured using Knox, and if you have upgraded from CDH 6 to CDP 7.1.7, then you may see a blank Hue screen. This could happen because the `knox_proxyhosts` parameter is newly introduced in CDP, and it is possible that this parameter is not configured in Cloudera Manager under Hue configuration.

Specify the host on which you have installed Knox in the Hue Knox Proxy Hosts configuration as follows:

1. Log in to Cloudera Manager as an Administrator.
2. Obtain the host name of the Knox Gateway by going to Clusters Knox service Instances .
3. Go to Clusters Hue service Configuration and search for the Knox Proxy Hosts field.



**Note:** Cloudera Manager displays the following warning if the Knox Proxy Hosts field is empty when Knox Gateway is enabled on the CDP cluster: The parameter `knox_proxyhosts` cannot be empty. This can happen if there are no Knox Gateways. Please set the `knox_proxyhosts` to the list of hosts that have Knox Gateways.

4. Specify the Knox Gateway hostname in the Knox Proxy Hosts field.
5. Click Save Changes and restart the Hue service.

#### **OPSAPS-58927: Connection failed error when accessing the Search app (Solr) from Hue**

If you are using Solr with Hue to generate interactive dashboards and for indexing data, and if you have deployed two Solr services on your cluster and selected the second one as a dependency for Hue, then Cloudera Manager assigns the hostname of the first Solr service and the port number of the second Solr service generating an incorrect Solr URL in the search section of the `hue.ini` file. As a result, you may see a “Connection failed” error when you try to access the Search app from the Hue web UI.

1. Log into Cloudera Manager as an Administrator.
2. Go to Clusters Hue service Configuration and add the following lines in the Hue Service Advanced Configuration Snippet (Safety Valve) for `hue_safety_valve.ini` field:

```
[search]
URL of the Solr Server
solr_url=http://[***HOSTNAME***]:[***PORT***]/solr/
```

For example:

```
solr_url=http://solr2:4567/solr/
```

3. Click Save Changes.
4. Restart the Hue service.

**CLR-72251: Invalid S3 URI error while accessing S3 bucket**

The Hue Load Balancer merges the double slashes (//) in the S3 URI into a single slash (/) so that the URI prefix `"/filebrowser/view=S3A:/"` is changed to `"/filebrowser/view=S3A:/"`. This results in an error when you try to access the S3 buckets from the Hue File Browser through the port 8889.

The Hue web UI displays the following error: “Unknown error occurred”.

The Hue server logs record the “ValueError: Invalid S3 URI: S3A” error.

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.

You should be able to load the S3 browser from both 8888 and 8889 ports.

Alternatively, you can use the Hue server port 8888 instead of the load balancer port 8889 to resolve this issue.

**CLR-72255: Error while rerunning Oozie workflow**

You may see an error such as the following while rerunning an already executed and finished Oozie workflow through the Hue web interface: `E0504: App directory [hdfs://cdh/user/hue/oozie/workspaces/hue-oozie-1571929263.84] does not exist`.

To resolve this issue, add the following property in the Hue Load Balancer Advanced Configuration Snippet:

1. Sign in to Cloudera Manager as an administrator.
2. Go to `Clusters Hue service Configurations Load Balancer` and search for the Load Balancer Advanced Configuration Snippet (Safety Valve) for `httpd.conf` field.
3. Specify `MergeSlashes OFF` in the Load Balancer Advanced Configuration Snippet (Safety Valve) for `httpd.conf` field.
4. Click Save Changes.
5. Restart the Hue Load Balancer.

**CDPD-16407: Python-psycopg2 package version 2.8.4 not compatible with Hue**

Ubuntu 18.04 provides `python-psycopg2` package version 2.8.4 but it is not compatible with Hue because of a bug in the Django framework.

Downgrade the package at the OS level by running the following command:

```
sudo apt install python-psycopg2==2.7.5
```

or install `python-psycopg2` package using pip by running the following command:

```
sudo pip install psycopg2==2.7.5
```

**DOCS-6344: Hue limitation after upgrading from CDH to CDP Private Cloud Base**

The `hive.server2.parallel.ops.in.session` configuration property changes from `TRUE` to `FALSE` after upgrading from CDH to CDP Private Cloud Base. Current versions of Hue are compatible with

this property change; however, if you still would like to use an earlier version of Hue that was not compatible with this property being FALSE and shared a single JDBC connection to issue queries concurrently, the connection will no longer work after upgrading.

**CDPD-43956: Manually replace UUID when importing Oozie workflows containing sub-workflows**

When importing Oozie workflows that contain sub-workflows, you must replace all the UUID entries with unique new entries. If you change the UUID of a sub-workflow, you must update that reference in the parent workflow to avoid circular dependencies.

**INSIGHT-3707: Query history displays "Result Expired" message**

You see the "Result Expired" message under the Query History column on the **Queries** tab for queries which were run back to back. This is a known behaviour.

None.

**Unsupported features****CDPD-59595: Spark SQL does not work with all Livy servers that are configured for High Availability**

SparkSQL support in Hue with Livy servers in HA mode is not supported. Hue does not automatically connect to one of the Livy servers. You must specify the Livy server in the Hue Advanced Configuration Snippet as follows:

```
[desktop]
[spark]
livy_server_url=http(s)://[***LIVY-FOR-SPARK3-SERVER-HOST***]:
[***LIVY-FOR-SPARK3-SERVER-PORT***]
```

Moreover, you may see the following error in Hue when you submit a SparkSQL query: Expecting value: line 2 column 1 (char 1). This happens when the Livy server does not respond to the request from Hue.

Specify all different Livy servers in the `livy_server_url` property one at a time and use the one which does not cause the issue.

**CDPD-18491: PySpark and SparkSQL are not supported with Livy in Hue**

Hue does not support configuring and using PySpark and SparkSQL with Livy in CDP Private Cloud Base.

**Importing and exporting Oozie workflows across clusters and between different CDH versions is not supported**

You can export Oozie workflows, schedules, and bundles from Hue and import them only within the same cluster if the cluster is unchanged. You can migrate bundle and coordinator jobs with their workflows only if their arguments have not changed between the old and the new cluster. For example, hostnames, NameNode, Resource Manager names, YARN queue names, and all the other parameters defined in the `workflow.xml` and `job.properties` files.

Using the import-export feature to migrate data between clusters is not recommended. To migrate data between different versions of CDH, for example, from CDH 5 to CDP 7, you must take the dump of the Hue database on the old cluster, restore it on the new cluster, and set up the database in the new environment. Also, the authentication method on the old and the new cluster should be the same because the Oozie workflows are tied to a user ID, and the exact user ID needs to be present in the new environment so that when a user logs into Hue, they can access their respective workflows.



**Note:** Migrating Oozie workflows from HDP clusters is not supported.

**Technical Service Bulletins****TSB 2023-704: File corruption when downloading files larger than 1 MB from ABFS with Hue File Browser**

An issue within the upstream Apache Hue (Hue) application results in file corruption when downloading files larger than 1MB files from Azure Blob Filesystem (ABFS) using the Hue File Browser. Only the downloaded files are affected by this issue, the source files remain intact.

#### Knowledge article

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2023-704: File corruption when downloading files larger than 1 MB from ABFS with Hue File Browser](#)

#### TSB 2023-703: Risk of Data Loss when using Hue S3 File Browser

There is a risk of losing data when moving one or more files or folders in Amazon S3 storage with the Hue File Browser. When the user selects the destination folder in the modal window the following scenarios can occur:

1. If the user selects the same destination folder as the source folder, and clicks the Move button, the selected files will be permanently deleted.
2. If the user selects a different destination folder from the source folder and clicks the Move button before the User Interface (UI) has completely loaded (the loading is indicated by a spinner), the action could lead to the following outcomes:
  - a. If the previously visible destination folder was the same as the source folder, the file will be permanently deleted.
  - b. If the previously visible destination folder was different from the source folder, the file(s) will be moved to the previously visible destination folder and not to the intended destination folder.

#### Knowledge article

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2023-703: Risk of Data Loss when using Hue S3 File Browser](#)

## Known Issues in Apache Iceberg

Learn about the known issues in Iceberg, the impact or changes to the functionality, and the workaround.

#### CDPD-75422: Impala schema case sensitivity issue with workaround

Impala's schema is case insensitive, causing errors with mixed case schema elements created through Spark during predicate pushdown.

- Create tables through Impala to ensure lower case schema.
- Avoid upper case in Spark: Do not use upper case when creating tables through Spark.
- Fix existing tables: Use ALTER TABLE to rename upper case columns:

```
ALTER TABLE `database`.`iceberg_table` CHANGE COLUMN ID id string;
```

#### 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

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.

## Technical Service Bulletins

#### TSB 2024-752: Dangling delete issue in Spark `rewrite_data_files` procedure causes incorrect results for Iceberg V2 tables

The Spark Iceberg library includes two procedures - `rewrite_data_files` and `rewrite_position_delete_files`. The current implementation of `rewrite_data_files` has a limitation that the position delete files are not deleted and still tracked by the table metadata, even if they

no longer refer to an active data file. This is called the dangling delete problem. To solve this, the `rewrite_position_delete_files` procedure is implemented in the Spark Iceberg library to remove these old “dangling” position delete files.

Due to the dangling delete limitation, when an Iceberg table with dangling deletes is queried in Impala, Impala tries to optimize `select count(*) from iceberg_table` query to return the results using stats. This optimization returns incorrect results.

The following conditions must be met for this issue to occur:

- All delete files in the Iceberg table are “dangling”
  - This would occur immediately after running `Spark rewrite_data_files` AND
    - Before any further delete operations are performed on the table OR
    - Before `Spark rewrite_position_delete_files` is run on the table
- Only stats optimized plain `select count(*) from iceberg_table` queries are affected. For example, the query should not have:
  - Any WHERE clause
  - Any GROUP BY clause
  - Any HAVING clause

Remove dangling deletes: After `rewrite_data_files`, position delete records pointing to the rewritten data files are not always marked for removal, and can remain tracked by the live snapshot metadata of the table. This is known as the dangling delete problem.

#### Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2024-752: Dangling delete issue in Spark `rewrite\_data\_files` procedure causes incorrect results for Iceberg V2 tables](#).

## Known Issues in Apache Impala

Learn about the known issues in Impala, the impact or changes to the functionality, and the workaround.

### CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

### CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the `--tls_ciphersuites` startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.



Ensure that the list in the `--tls_ciphersuites` startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

Apache Jira: [IMPALA-14625](#)

**CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.**

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

**CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work**

If you try to run `impala-shell` on RHEL 9 by setting the default python executable available in `PATH` to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, `IMPALA_PYTHON_EXECUTABLE=python3`.

**CDPD-80019: Impala fails to start due to large catalog-update topic**

Impala coordinators fail to start if the catalog-update statestore topic exceeds 2GB, hitting the Thrift 0.16.0 message size limit (`thrift_rpc_max_message_size`). This prevents the coordinator from processing the update.

- Enable local catalog mode: `--use_local_catalog=true`
- In Cloudera Manager, go to Impala Configuration search for "Catalog Server Command Line

Argument Advanced Configuration Snippet (Safety Valve)" and click **+**

- Add property name and value:

```
Name : --catalog_topic_mode
Value: minimal
```

- Save changes.

Apache Jira: [IMPALA-12350](#)

**CDPD-42958: After upgrading the CDH 7.1.9 from CDH 6.x, under certain conditions you cannot insert data into a table**

Under the following conditions, after upgrading from CDH 6.x to CDH 7.1.9 you cannot insert data into a table from Impala:

- On CDH 6.x, you created a database with Impala in a user specified HDFS location.
- Using Hive, you then created a table in the database.

Under these conditions, the database and table are stored in the user-specified HDFS directory. After upgrading, the HDFS directory of the table is read-only for Impala. Consequently, from Impala you cannot insert new data into the table because Impala does not have write permission on the HDFS directory.

Workaround: To resolve this issue, use either one of the following workarounds:

- Using the Ranger Web UI, in the policy repository `cm_hdfs`, grant the user 'impala' write permission on the directory where the table resides.
- Enter the following command to grant write permission to user 'impala' on the HDFS directory where the table resides.

```
hdfs dfs -setfacl -m default:user:impala:rwX <HDFS directory>
```

**Impala cannot update table if the 'external.table.purge' property is not set to true**

Impala cannot update a table using DDL statements if the 'external.table.purge' property is `FALSE`. `ALTER TABLE` statements return success with no changes to the table.

`ALTER TABLE` statements should be issued twice if "external.table.purge" was `FALSE` initially.



**Note:** For Iceberg tables, Impala users should always use CREATE TABLE since the Impala CREATE TABLE statement sets the flag to TRUE. However, Impala CREATE EXTERNAL TABLE sets the flag to FALSE.

### Impala's known limitation when querying compacted tables

When the compaction process deletes the files for a table from the underlying HDFS location, the Impala service does not detect the changes as the compactions does not allocate new write ids. When the same table is queried from Impala it throws a 'File does not exist' exception that looks something like this:

```
Query Status: Disk I/O error on <node>:22000: Failed to open HDF
S file hdfs://nameservice1/warehouse/tablespace/managed/hive/<da
tabase>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException:
File does not exist: /warehouse/tablespace/managed/hive/<data
base>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

### CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

You must use a single Knox node to access Impala UI.

### Impala api calls via kinox require configuration if the kinox customized kerberos principal name is a default service user name

To access impala api calls via kinox, if the kinox customized kerberos principal name is a default service user name, then configure "authorized\_proxy\_user\_config" by clicking Clusters->impala->configuration. Include the kinox customized kerberos principal name in the comma separated list of values <knox\_custom\_kerberos\_principal\_name>="\*" where <knox\_custom\_kerberos\_principal\_name> is the value of the Kerberos Principal in the Knox service. Select Clusters>Knox>Configuration and search for Kerberos Principal to display this value.

### CDPD-21828: Multiple permission assignment through grant is not working

None

### Problem configuring masking on tables using Ranger

The following Knowledge Base article describes the behavior when we configure masking on tables using Ranger. This configuration works for Hive, but breaks queries in some scenarios for Impala.

For a workaround, see the following Knowledge Base article: [ERROR: "AnalysisException: No matching function with signature: mask\(FLOAT\)" when Impala jobs fail with the following error with signature: mask\(FLOAT\)](#)

### IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm\_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXXXX"

### IMPALA-532: Impala should tolerate bad locale settings

If the LC\_\* environment variables specify an unsupported locale, Impala does not start.

Add LC\_ALL="C" to the environment settings for both the Impala daemon and the Statestore daemon.

### Avro Scanner fails to parse some schemas



```
| Explain String |
+-----+
- - - - - +
| Estimated Per-Host Requirements: Memory=1.00KB VCores=1 |
| 00:EMPTYSET |
+-----+
- - - - - +
```

### IMPALA-3509: Breakpad minidumps can be very large when the thread count is high

The size of the breakpad minidump files grows linearly with the number of threads. By default, each thread adds 8 KB to the minidump size. Minidump files could consume significant disk space when the daemons have a high number of threads.

Add `-minidump_size_limit_hint_kb=size` to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then 2 KB per thread after that.) The minidump file can still grow larger than the "hinted" size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

## IMPALA-4978: Impala requires FQDN from hostname command on Kerberized clusters

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the `gethostname()` system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a Kerberized cluster.

Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname`, only returns the short name, pass the command-line flag `##hostname=FULLY_QUALIFIED_DOMAIN_NAME` in the startup options of all Impala-related daemons.

## IMPALA-6671: Metadata operations block read-only operations on unrelated tables

Metadata operations that change the state of a table, like `COMPUTE STATS` or `ALTER RECOVER PARTITIONS`, may delay metadata propagation of unrelated unloaded tables triggered by statements like `DESCRIBE` or `SELECT` queries.

None

## IMPALA-7072: Impala does not support Heimdal Kerberos

None

**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.

## 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.

Some of the unresolved issues include:

- IMPALA-6841
- IMPALA-635

## Known Issues in Apache Kafka

Learn about the known issues in Kafka, the impact or changes to the functionality, and the workaround.

### CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

### 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-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL\_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker) in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Save Changes > Restart SMM.

### The offsets.topic.replication.factor property must be less than or equal to the number of live brokers

The offsets.topic.replication.factor broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a GROUP\_COORDINATOR\_NOT\_AVAILABLE error until the cluster size meets this replication factor requirement.

None

### Requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true

The first few produce requests fail when sending to a nonexistent topic with auto.create.topics.enable set to true.

Increase the number of retries in the producer configuration setting retries.

**Performance degradation when SSL Is enabled**

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with `ssl.secure.random.implementation = SHA1PRNG`. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

**OPSAPS-43236: Kafka garbage collection logs are written to the process directory**

By default Kafka garbage collection logs are written to the agent process directory. Changing the default path for these log files is currently unsupported.

None

**RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure**

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transactional method because we are in an error state
 at org.apache.kafka.clients.producer.internals.TransactionManager.maybeFailWithError(TransactionManager.java:1125)
 at org.apache.kafka.clients.producer.internals.TransactionManager.maybeAddPartition(TransactionManager.java:442)
 at org.apache.kafka.clients.producer.KafkaProducer.doSend(KafkaProducer.java:1000)
 at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:914)
 at org.apache.kafka.clients.producer.KafkaProducer.send(KafkaProducer.java:800)
 .
 .
 .
Caused by: org.apache.kafka.common.errors.ClusterAuthorizationException: Cluster authorization failed.
```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting `enable.idempotence` to `false`.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

**CDPD-49304: AvroConverter does not support composite default values**

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

**DBZ-4990: The Debezium Db2 Source connector does not support schema evolution**

The Debezium Db2 Source connector does not support the evolution (updates) of schemas. In addition, schema change events are not emitted to the schema change topic if there is a change in the schema of a table that is in capture mode. For more information, see [DBZ-4990](#).

None.

**CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression**

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy

compression. The HDFS Stateless Sink connector is only affected if the Compression Codec or Compression Codec for Parquet properties are set to SNAPPY.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: org.apache.hadoop.util.NativeCodeLoader.buildSupportsSnappy()
```

```
Failed to write to HDFS due to java.lang.RuntimeException: native snappy library not available: this version of libhadoop was built without snappy support.
```

Download and deploy missing libraries.



**Important:** Ensure that you complete steps 1-11 on all Kafka Connect hosts. Additionally, ensure that the advanced configuration snippet in step 12 is configured for all Kafka Connect role instances.

1. Create the /opt/nativelibs directory.

```
mkdir /opt/nativelibs
```

2. Change the owner to kafka.

```
chown kafka:kafka /opt/nativelibs
```

3. Locate the directory containing the Hadoop native libraries and copy its contents to the directory you created.

```
cp /opt/cloudera/parcels/CDH/lib/hadoop/lib/native/* /opt/nativelibs
```

4. Verify that 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-69481: Some Kafka Connect metrics missing from Cloudera Manager due to conflicting definitions**

The metric definitions for kafka\_connect\_connector\_task\_metrics\_batch\_size\_avg and kafka\_connect\_connector\_task\_metrics\_batch\_size\_max in recent Kafka CSDs conflict with previous definitions in other CSDs. This prevents Cloudera Manager from registering these metrics. It also results in SMM returning an error. The metrics also cannot be monitored in Cloudera Manager chart builder or queried using the Cloudera Manager API.

Contact Cloudera support for a workaround.

#### **OPSAPS-71258: Zstandard and Snappy compression do not support /tmp mounted as noexec**

Kafka cannot process messages compressed with Zstandard or Snappy if /tmp is mounted as noexec.

The workaround steps for Zstandard compression are the following. You need to complete these steps for all the collected hosts.

1. Identify the hosts running Kafka.
2. Verify the Zstandard version by checking the classpath of the running process.

```
ps aux
```

By default, Kafka uses version libzstd-jni-1.5.2-1.

The Zstandard JAR file is located at /opt/cloudera/parcels/CDH/jars/zstd-jni-1.5.2-1.jar. This file includes the architecture and OS-specific binaries. Check [Zstandard JNI GitHub repository](#) for the full list of supported native libraries.

3. Extract the binary corresponding to the target system similar to the following example.

```
/usr/java/default/bin/jar xf /opt/cloudera/parcels/CDH/jars/zstd-jni-1.5.2-1.jar linux/amd64/libzstd-jni-1.5.2-1.so
```

4. Copy the extracted binary to a location included in java.library.path similar to the following example.

```
cp linux/amd64/libzstd-jni-1.5.2-1.so /lib
```

5. Restart Kafka.

The workaround steps for Snappy compression are the following:

1. In Cloudera Manager, select the Kafka service and go to Configuration.



2. Add the following line to Additional Broker Java Options.

```
-org.xerial.snappy.tmpdir==[***PATH***]
```

Where [\*\*\*PATH\*\*\*] is a directory that is not /tmp.

3. Restart Kafka.

## Unsupported Features

The following Kafka features are not supported in Cloudera Data Platform:

- Only Java and .Net based clients are supported. Clients developed with C, C++, Python, and other languages are currently not supported.
- The Kafka default authorizer is not supported. This includes setting ACLs and all related APIs, broker functionality, and command-line tools.
- SASL/SCRAM is only supported for delegation token based authentication. It is not supported as a standalone authentication mechanism.
- Kafka KRaft in this release of Cloudera Runtime is in technical preview and does not support the following:
  - Deployments with multiple log directories. This includes deployments that use JBOD for storage.
  - Delegation token based authentication.
  - Migrating an already running Kafka service from ZooKeeper to KRaft.
  - Atlas Integration.

## Limitations

### Collection of Partition Level Metrics May Cause Cloudera Manager's Performance to Degrade

If the Kafka service operates with a large number of partitions, collection of partition level metrics may cause Cloudera Manager's performance to degrade.

If you are observing performance degradation and your cluster is operating with a high number of partitions, you can choose to disable the collection of partition level metrics.



**Important:** If you are using SMM to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both SMM and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, SMM will not be able to display information about partitions. In addition, Cruise Control will not operate properly.

Complete the following steps to turn off the collection of partition level metrics:

1. Obtain the Kafka service name:
  - a. In Cloudera Manager, Select the Kafka service.
  - b. Select any available chart, and select Open in Chart Builder from the configuration icon drop-down.
  - c. Find \$SERVICENAME= near the top of the display.

The Kafka service name is the value of \$SERVICENAME.

2. Turn off the collection of partition level metrics:
  - a. Go to Hosts Configuration.
  - b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

Enter the following to turn off the collection of partition level metrics:

```
[KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_entity_update_enabled=false
```

Replace [KAFKA\_SERVICE\_NAME] with the service name of Kafka obtained in step 1. The service name should always be in lower case.

- c. Click Save Changes.

## Known Issues in Kerberos

Learn about the known issues in Kerberos, the impact or changes to the functionality, and the workaround.

## Key Trustee Server

Learn about the known issues in Key Trustee Server, the impact or changes to the functionality, and the workaround.

### KTS setup on FIPS cluster gets stuck if KTS is in HA

Key Trustee Server setup gets stuck if KTS is in HA for Centos 7.9 in FIPS mode.

### Create Date is not fetched from Key Trustee Server

During key migration from KTS to Ranger KMS , create date is not fetched from KTS.

## Known Issues in Apache Knox

Learn about the known issues in Knox, the impact or changes to the functionality, and the workaround.

### CDPD-61088: When downgrade is performed from CDP 7.1.9 to CDP 7.1.7 SP2, Knox may fail to start.

Failed to start gateway: org.apache.knox.gateway.services.ServiceLifecycleException: Keystore was not loaded properly - the provided password may not match the password for the keystore. org.apache.knox.gateway.services.ServiceLifecycleException: Keystore was not loaded properly - the provided password may not match the password for the keystore.

Workaround: Remove the faulty credential store and restart Knox.

### CDPD-60996: When downgrade is performed from CDP 7.1.9 to CDP 7.1.7 SP2, Knox is unable to connect to Cloudera Manager.

Restart Knox service after downgrade.

### CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

You must use a single Knox node to access Impala UI.

### CDPD-3125: Logging out of Atlas does not manage the external authentication

At this time, Atlas does not communicate a log-out event with the external authentication management, Apache Knox. When you log out of Atlas, you can still open the instance of Atlas from the same web browser without re-authentication.

To prevent additional access to Atlas, close all browser windows and exit the browser.

### CDPD-22785: Improvements and issues needs to be addressed in convert-topology knox cli command

None

**OPSAPS-67480:** In 7.1.9, default Ranger policy is added from the cdp-proxy-token topology, so that after a new installation of CDP-7.1.9, the Knox-Ranger policy includes cdp-proxy-token. However, upgrades do not add cdp-proxy-token to cm\_knox policies automatically.

Manually add cdp-proxy-token to the Knox policy, using Ranger Admin Web UI.

1. Log in to Cloudera Manager Ranger Admin Web UI, as a Ranger administrator.
2. On Ranger Admin Web UI Service Manager Resource Knox, click cm\_knox.
3. In Knox Policies, open the CDP Proxy UI, API and Token policy.
4. In Knox Topology\*, add cdp-proxy-token.
5. Click Save.
6. Restart Ranger.

## Known Issues in Apache Kudu

Learn about the known issues in Kudu, the impact or changes to the functionality, and the workaround.

- Kudu HMS Sync is disabled and is not yet supported

### You get "The user 'kudu' is not part of group 'hive' on the following hosts: " warning by the Host Inspector

If you are using fine grained authorization for Kudu, and you are also using Kudu-HMS integration with HDFS-Sentry sync, then you may get the "The user 'kudu' is not part of group 'hive' on the following hosts: " warning while upgrading.

Workaround: Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

## Known Issues in Navigator Encrypt

Learn about the known issues in Navigator Encrypt, the impact or changes to the functionality, and the workaround.

### Failed to start navencrypt-mount.service

After installing NavEncrypt on SLES, if the command "systemctl status navencrypt-mount" fails with the error: "Failed to start navencrypt-mount.service: Unit navencrypt-mount.service failed to load: No such file or directory", the systemd files need to be reloaded.

Run the following command:

```
(sudo) systemctl daemon-reload
```

### Page allocation failure

If a customer is running RHEL 8.8 (which is Linux Kernel 4.18), it is possible for the kernel to report a page allocation failure and cause a kernel panic.

Delete all ACLs that Navigator Encrypt is running with.

### Crash in Navigator Encrypt kernel module if --profile-file= option is used with ACLs

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.

### OPSAPS-74399: Missing jars from Spark 3 ShareLib when using with CDS 3.5

The spark-common-utils\_2.12-3.5.4.3.5.7191000.0-XX.jar and spark-sql-api\_2.12-3.5.4.3.5.7191000.0-XX.jar files are missing from the Spark 3.5.4 sharelib directory causing the Spark 3 actions to fail.

Perform the troubleshooting steps as described in [Troubleshoot Spark 3 ShareLib compatibility behaviour](#).

#### CDPD-41274: HWC + Oozie issue: Could not open client transport with JDBC Uri

Currently only Spark cluster mode is supported in the Oozie Spark Action with Hive Warehouse Connector (HWC).

Use Spark action in cluster mode.

```
Use Spark action in cluster mode.
 <spark xmlns="uri:oozie:spark-action:1
.0">
 ...
 <mode>cluster</mode>
 ...
</spark>
```

#### CDPD-26975: Using the ABFS / S3A connectors in an Oozie workflow where the operations are "secured" may trigger an IllegalArgumentException with the error message java.net.URISyntaxException: Relative path in absolute URI.

Set the following XML configuration in the Datahub cluster's Cloudera Manager:

1. In the Cloudera Manager Admin Console, go to the Oozie service.
2. Click the Configuration tab.
3. In the Oozie Server Advanced Configuration Snippet (Safety Valve) for oozie-site.xml field, set the following:

Set the following if you are using Amazon S3:

```
<property>
<name>oozie.service.HadoopAccessorService.fs.s3a</name>
<value>fs.s3a.buffer.dir=/tmp/s3a</value>
</property>
```

Set the following if you are using ABFS:

```
<property>
<name>oozie.service.HadoopAccessorService.fs.abfs</name>
<value>fs.azure.buffer.dir=/tmp/abfs</value>
</property>

<property><name>
oozie.service.HadoopAccessorService.fs.abfss</name>
<value>fs.azure.buffer.dir=/tmp/abfss</value>
</property>
```

4. Enter a Reason for change, and then click Save Change to commit the changes.
5. Restart the Oozie service.

#### OOZIE-3549: Oozie fails to start when Cloudera Manager 7.x is used with Cloudera Runtime 6.x and Java 11 because Oozie does not set the trust-store password.

The issue is fixed in OOZIE-3549 and is already included in CDP 7.x but not in CDH 6.x. If you are on CDH 6.x and want to upgrade to Java 11 or your Cloudera Manager to 7.x then you must request for a patch.

#### 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.

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.

## Unsupported Feature

The following Oozie features are currently not supported in Cloudera Data Platform:

- Non-support for Pig action (CDPD-1070)
- Conditional coordinator input logic

Cloudera does not support using Derby database with Oozie. You can use it for testing or debugging purposes, but Cloudera does not recommend using it in production environments. This could cause failures while upgrading from CDH to CDP.

## Known Issues in Apache Ozone

Learn about the known issues in Ozone, the impact or changes to the functionality, and the workaround.

### **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: 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.

### **The Prometheus binary is not available in CDP Private Cloud Base 7.1.9 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.

### **After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and if Ozone is in the Non-HA environment, an exception message is observed during the finalization of the Ozone upgrade.**

During the finalization of the upgrade, a `ClassNotFoundException` message for `org.cloudera.log4j.redactor.RedactorAppender` class may be displayed. The error message can be ignored because the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

### **CDPD-68951: In CDP Private Cloud Base 7.1.9 CHF2 version and lower, the command `ozone sh key list <bucket_path>` displays the `isFile` flag in a key's metadata as `false` even when the key is a file. This issue is rectified in CDP Private Cloud Base 7.1.9 CHF3. However, the pre-existing (pre-upgrade) key's metadata cannot be changed.**

None

**CDPD-67771: When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.**

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(OzoneManager
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.lambda
$1(OzoneManagerStateMachine.java:363) at
java.base/java.util.concurrent.CompletableFuture
$AsyncSupply.run(CompletableFuture.java:1700) at java.base/
java.util.concurrent.ThreadPoolExecutor.runWorker(ThreadPoolExecutor.java:1128)
at java.base/java.util.concurrent.ThreadPoolExecutor
$Worker.run(ThreadPoolExecutor.java:628) at java.base/
java.lang.Thread.run(Thread.java:834)
```

This occurs because S3A uses multipart upload to commit job results in a batch. The staging committer's replace mode deletes the target directory before completing Memory Protection Unit (MPU). The problem is that File System Optimization (FSO) does not create intermediate directories during MPU, it does only for regular file/dir/key requests.

Use `fs.s3a.committer.name=magic` for affected versions.

**CDPD-64398: In SCM with non-HA configuration, Secret key manager does not initialise. Hence, the startup of OM and DataNode fails because it cannot get a secret key. This key is used when security is enabled and used in Block token and container token verification while communicating between the Ozone client, OM, and Datanode.**

This issue does not occur in SCM with HA configuration.

You must force exit from the safe mode for SCM. This triggers the initialization of the Secret key manager. Below are the options:

- exit safe mode manually whenever SCM is started
- disable safe mode by setting `hdds.scm.safemode.enabled=false` in safety valve for SCM configuration

Impact of disabling safe mode and the purpose of safe mode:

- Write should not fail once the cluster is out of safe mode
- Read of existing data should not fail after the cluster is out of safe mode
- Unnecessary re-replication should be avoided during cluster restart

So disabling safe mode does not have any major impact on Ozone.

**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.

**CDPD-52412: keyManagerImpl#listStatus exceeds the maximum RPC length.**

To reduce the payload, you must increase the part size.

**OPSAPS-68159: If you did not deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Error when distributing parcel to host, Ignoring non-compliant parcel manifest error is displayed after Cloudera Manager is upgraded to 7.11.3.**

If you encounter the error, perform the following steps:

1. Deactivate and undistribute the Ozone parcel 718.1.0 on Cloudera Manager 7.11.3.
2. Restart the cluster with a delay of 10 minutes.
3. Continue to upgrade CDH 7.1.8 to CDH 7.1.9.

**OPSAPS-68159: If you did not deactivate the Ozone parcel 718.2.x on Cloudera Manager 7.7.1 with CDH 7.1.8 before upgrading to Cloudera Manager 7.11.3 with CDH 7.1.9, the Ozone roles during the CDH 7.1.8 upgrade to CDH 7.1.9.**

If you encounter the error, perform the following steps:

1. Deactivate the Ozone parcel 718.2.x.
2. Restart the Ozone service.
3. Perform Finalize Upgrade for the Ozone service.

Step result: The Ozone roles will be displayed in green.

**CDPD-60989: The packaging version for Apache Ozone points to the 1.2.0 older version. This is a version string problem and not a packaging issue. The version of the Apache Ozone binary is closest to 1.4.0.**

None. This only affects the JAR versioning.

**CDPD-60366: Native library loader fails when system property native.lib.tmp.dir is not set. It fails because the library is copied to / instead of the cwd.**

Setting native.lib.tmp.dir to a certain path like /tmp should solve the issue. To set this system property, you must add -Dnative.lib.tmp.dir=/tmp to ozone\_java\_opts in Cloudera Manager configuration of the Ozone cluster.

**CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.**

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

**CDPD-60598: DataNode can consume up to the JVM heap size worth of direct memory buffers under specific failure scenarios. This can lead to a DataNode crash. This is due to Netty allocating and not freeing the direct memory buffers.**

Restarting DataNode will resolve this issue.

**CDPD-60466: The Ozone client is missing in the Cloudera Manager's Service Monitor node. This causes Ozone Canary's health check to fail.**

Ensure you install Ozone CLI on the same node as Cloudera Manager's Service Monitor.

**CDPD-60012: The Ozone SCM may be stuck in safe mode after upgrade to 7.1.9. This is due to a bug in how the SCM accounts for open storage containers.**

The workaround is to temporarily set the Ozone configuration hdds.scm.safemode.threshold.pct to a lower value like 0.90 and restart the SCM.

**CDPD-59679: In some scenarios, excess UNHEALTHY replicas of EC containers may not be removed from the cluster.**

None

**CDPD-50116: Topology-aware reads can provide better performance by routing applications to the closest replica of a file, if available, on the same physical node or same rack. However this feature is disabled by default in CDP 7.1.9.**

This feature can be enabled by setting ozone.network.topology.aware.read to true using Cloudera Manager safety valve.

**CDPD-57165: The DataNode disk usage thread may abort due to Ratis log tailing.**

Use `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory` to calculate disk usage. Configure in the `ozone-site.xml` property: `hdds.datanode.du.factory` value: `org.apache.hadoop.hdds.fs.DedicatedDiskSpaceUsageFactory`

**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.

**CDPD-55513: Hive external table replication policy does not migrate Hive tables in HDFS storage to Ozone.**

To replicate data from HDFS to Ozone, see [Migrating your data from HDFS to Ozone](#). You can use HMS Mirror to replicate the metadata.

**OPSAPS-67757: Hive external tables in Ozone storage cannot be replicated using Hive external table replication policies.**

To replicate the Hive external tables' data, consider using DistCp. To replicate the metadata of Hive external tables, consider using HMS Mirror.

**Remove bucket recursively using `rb --force` command from AWS S3 cannot work for FSO buckets.**

Use the Ozone shell command `ozone sh bucket delete -r [***BUCKET ADDRESS***]`

**CDPD-59126: Info log displays "noexec permission on /tmp/liborg\_apache\_ratis\_thirdparty\_netty\_transport\_native\_epoll\_x86" on client while executing command with noexec on /tmp.**

To suppress Info log related to `liborg_apache_ratis_thirdparty_netty_transport_native_epoll_x86` library: Export `OZONE_OPTS` environment variable on the client terminal by running the command `export OZONE_OPTS="-Dorg.apache.ratis.thirdparty.io.netty.native.workdir=/var/tmp $OZONE_OPTS"`

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-64097: Ozone service restart failed at SCM**

Stopping SCM service using Cloudera Manager can sometimes timeout and need a retry. The Cloudera Manager API waits for 90 seconds which is not sufficient under certain circumstances.

Retry the shutdown using Cloudera Manager if the SCM still shows as running after a refresh of the service status.

**OPSAPS-66469: Ozone-site.xml is missing if the host does not contain HDFS roles**

The client side `ozone-site.xml` (`/etc/hadoop/conf/ozone-site.xml`) is not generated by Cloudera Manager if the host does not have any HDFS role. Because of this, issuing Ozone commands from that host fails because it cannot find the service name to host name mapping. The error message is similar to this: `# ozone sh volume list o3://ozoneabc 23/03/06 18:46:15 WARN ha.OMProxyInfo: OzoneManager address ozoneabc:9862 for serviceID null remains unresolved for node ID null` Check your `ozone-site.xml` file to ensure ozone manager addresses are configured properly.

Add the HDFS gateway role on that host.

**OPSAPS-67607: Cloudera Manager FirstRun failure at the “Upload YARN MapReduce Framework JARs” step.**

If this failure is attributed to the broken symbolic link, `/var/lib/hadoop-hdfs/ozone-file-system-hadoop3.jar`, it is likely due to the presence of the user `hdfs` on the node prior to CDP parcel activation. As a result, the Cloudera Manager agent skips the initialization related to HDFS, leading to the non-creation of the `/var/lib/hadoop-hdfs` directory.

Create the directory `“/var/lib/hadoop-hdfs”` on all nodes followed by the deactivation and activation of the CDP parcel (deactivate and activate the Ozone parcel instead, in case Ozone parcel is used).

**CDPD-50447: When SCM High Availability is enabled, each of the SCM web UIs report the host of the web UI as the leader of HA, and the other two as followers. This gives wrong information**

Correct output is available by running the `ozone admin scm roles --service-id=<ID>` command.

**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](#).

**OPSAPS-62327: In an Ozone cluster without any gateway roles, Ozone is unable to deploy client configurations and displays the ConfigGenException error.**

You must add the Ozone gateway roles to the cluster.

**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-36389: The configurations `datanodes.involved.max.per.iteration` and `size.moved.max.per.iteration` are meant to limit the max number of DataNodes that is involved and maximum size that can move in an iteration. This bug causes balancer to stop an iteration when it is two DataNodes or 1 Container size (5GB) away from hitting these limits. However, these DataNodes can again be considered for balancing in the next iteration. This means the cluster is balanced after enough iterations, albeit a bit slowly. This bug is apparent in small clusters of around four DataNodes where the DataNodes can be either the source or target for a lot of moves but the iteration gets stopped when three DataNodes have been involved. It takes a higher number of iterations to eventually balance this cluster. While this is a performance issue, it does not prevent balancer from ultimately balancing the cluster.**

Increase the speed for balancing by decreasing the interval between each iteration using the configuration `balancing.iteration.interval`. The value of this configuration must be greater than `hdds.datanode.du.refresh.period`. `size.moved.max.per.iteration` can be increased to allow more data to move in one iteration.

**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-35141: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2> (state=08S01,code=40000) java.sql.SQLException: Error while compiling statement: FAILED: Execution Error, return code 40000 from org.apache.hadoop.hive ql.exec.MoveTask. Unable to move source <bucket1> to destination <bucket2>. This issue does not occur if the source and target buckets are different in Hive queries. Currently, copying across the same bucket is only supported.**

Avoid different buckets in source and target path. Ensure that `hive.exec.stagingdir` is unset or set to the same bucket as the warehouse so that staging directories are created in the same bucket as the warehouse.

**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-41184: With Legacy buckets, FileSystem Optimised is not interoperating with the Ozone shell command. The cause is that the directory key entry in the DB KeyTable is stored as "dir1/" with trailing slash. But while performing the described operation, Ozone shell (o3://) is normalizing the given path and removing the trailing slash "/" from it. That results in KEY\_NOT\_FOUND exception.**

There are three workarounds:

- Use FileSystem API to Delete the Directories rather than Shell-Command API.
- Use FSO buckets instead of Legacy Buckets. As in FSO, you can create Intermediate Directories and Delete Directories using the Ozone shell commands.
- Disable and set the configuration `ozone.om.enable.filesystem.pathsflag` to false in order to delete the directories. This is generally not a preferred workaround because the cluster must be restarted to pick up the new changes.

**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-12542: Mount of Ozone filesystem with the help of FUSE fails.**

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-42897: EC writes are failing with "No enough datanodes to choose" after EC replication config is set globally.**

EC writes start failing when large number of pipelines are created as a result of multiple EC configs with different chunk sizes used to write keys.

If standard EC configs (that is., rs-3-2-1024k) are used to write keys, number of pipelines created per DataNode is limited to five and this issue is not seen with standard EC configs.

The recommendation is not to create too many random chunk sizes. It is configurable because, users can decide based on their workload. But do not to have separate chunk sizes for each file.

**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-40560: Filesystem Operations using Hadoop S3a connector on a FILE\_SYSTEM\_OPTIMIZED bucket fails. `org.apache.hadoop.ozone.om.exceptions.OMException: Unable to get file status: volume: s3v bucket: fso key: test/`**

Do not run Hadoop S3a commands on a FILE\_SYSTEM\_OPTIMIZED bucket. Use OBJECT\_STORE bucket layouts.

**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.

**CDPD-43432: Ozone Service in fault state in DataNode - Long Running setup.**

Upgrade RocksDB to the latest version.

**OPSAPS-63999: In the newly installed cluster, the Finish upgrade option is clickable.**

None

**OPSAPS-64648: Failed to start Ozone node using Cloudera Manager if default log path /var/log/hadoop-ozone does not exist. If this path does not exist, any Ozone nodes(for example SCM or DataNode) restart will fail.**

Run the following command `sudo -u hdfs mkdir -p /var/log/hadoop-ozone` or replace hdfs with the user Ozone roles that are running.

**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.

**Technical Service Bulletins****TSB 2024-724: Issue with Apache Ozone non-HA SCMs during installation/upgrade when security tokens are enabled**

Apache Ozone (Ozone) Storage Container Manager (SCM) in a non-HA configuration with security tokens enabled may fail to start once the cluster is upgraded to or newly installed with Cloudera Data Platform (CDP) 7.1.9.

The Ozone Manager (OM) and Datanodes roles are unable to get secret keys from the SCM and are therefore unable to use them for token verification for containers and blocks. SCM contains a secret key manager service that generates and provides secret keys to OM and datanodes and the initialization does not occur post-upgrade.

**Knowledge article**

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2024-724: Issue with Apache Ozone non-HA SCMs during installation/upgrade when security tokens are enabled](#)

**TSB 2024-749: Possible Ozone Snapshot Chain Corruption in 7.1.9.0**

The Apache Ozone (Ozone) snapshots feature provides the ability to save consistent and immutable copies of the Ozone namespace. A chain of snapshots is tracked and maintained in the Ozone RocksDB that is utilized to calculate the difference between two snapshots required for the snapDiff API. A race condition on the snapshot delete path was found during testing. This condition might result in two consecutive snapshots pointing to the same parent, which breaks the snapshot chain.

**Upstream JIRA**

[HDDS-10524](#), [HDDS-10590](#), and [HDDS-9198](#)

**Knowledge article**

For the latest update on this issue see the corresponding Knowledge article: [TSB 2024-749: Possible Ozone Snapshot Chain Corruption in 7.1.9.0](#)

**TSB 2025-820: Potential Data Integrity Issues Found in Ozone**

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

#### Knowledge article

For the latest update on this issue see the corresponding Knowledge 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

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.

#### Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

### Known Issues in Apache Parquet

There are no known issues for Parquet in Cloudera Runtime 7.1.9.

### Known Issues in Apache Phoenix

Learn about the known issues in Phoenix, the impact or changes to the functionality, and the workaround.

**CDPD-60862: DDL refers to "Data Definition Language", a subset of SQL statements that change the structure of the database schema in some way, typically by creating, deleting, or modifying schema objects such as databases, tables, and views. Cloudera services that support DDL include Impala, Hive (using HiveQL), Spark (using SparkSQL), HBase, Phoenix, Flink, and Kafka.**

1. While ZDU upgrades are in flight, the CDP runtime is being upgraded from a lower version to a higher version in a rolling fashion.
2. It is mandatory that you do not run a DDL statement while the ZDU activity is in progress.
3. You can continue to use DML statements normally during ZDU.
4. After the ZDU has completed and the master and worker nodes have been restarted, you can resume running all statements including DDL and DML.

Since this is an application error, there is no workaround. Cloudera recommends you to modify the application to not use DDLs for the duration of the upgrade. After removing the DDLs, the application should work normally and proceed from the point of failure.

**CDPD-35925: Omid service fails to start rarely. Restarting it solves the problem most of the times.**

None

## Known Issues in Queue Manager

Learn about the known issues in Queue Manager, the impact or changes to the functionality, and the workaround.

**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.

## Known Issues in Apache Ranger

Learn about the known issues in Ranger, the impact or changes to the functionality, and the workaround.

**OPSAPS-75673: Wrong enablement of Ranger RMS Database Full Sync command**

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

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.

**CDPD-66092: Upgrade from CDP-7.1.8 to 7.1.9 fails, Ranger shows no policies at all, due to Java patches not being applied during upgrade. This causes rolling restart of services to fail.**

Skip applying any of the following java patches NOT applicable to the underlying environment. In other words, do not apply patches for a service definition that does not appear in ranger database.

**Table 8: Java Patches and Related Service Definitions**

| Java patch                                          | Service-definition |
|-----------------------------------------------------|--------------------|
| PatchForHiveServiceDefUpdate_J10027.java            | Hive               |
| PatchForTagServiceDefUpdate_J10028.java             | Tag                |
| PatchForHBaseServiceDefUpdate_J10035.java           | HBase              |
| PatchForHdfsAddChainedPluginProvider_J10038.java    | HDFS               |
| PatchForHdfsRemoveChainedPluginProvider_J10039.java | HDFS               |
| PatchForOzoneServiceDefUpdate_J10041.java           | Ozone              |
| PatchForOzoneServiceDefConfigUpdate_J10051.java     | Ozone              |
| PatchForOzoneServiceDefUpdate_J10057.java           | Ozone              |

**CDPD-61050: When there are no roles defined in Ranger, creating ranger replication policy causes empty export roles file to be created which causes transform step to fail.**

If source Ranger has at least 1 role, then the issue does not appear. For this, you can create a dummy role on the source Ranger.

**CDPD-60489: Jackson-dataformat-yaml 2.12.7 and Snakeyaml 2.0 are not compatible.**

You must not use Jackson-dataformat-yaml through the platform for YAML parsing.

**CDPD-56803: When there is no existing policy for user and a revoke request comes from hbase, then will get this error.**

```
hbase:001:0> revoke 'hrt_11'
ERROR: org.apache.hadoop.hbase.coprocessor.CoprocessorException:
HTTP 400 Error: processSecureRevokeRequest processing failed
at org.apache.ranger.authorization.hbase.RangerAuthorizationCopr
ocessor.preRevoke(RangerAuthorizationCoprocessor.java:1309)
at org.apache.ranger.authorization.hbase.RangerAuthorizationCop
rocessor.preRevoke(RangerAuthorizationCoprocessor.java:1128)
at org.apache.hadoop.hbase.master.MasterCoprocessorHost$162.ca
ll(MasterCoprocessorHost.java:1857)
at org.apache.hadoop.hbase.master.MasterCoprocessorHost$162.cal
l(MasterCoprocessorHost.java:1854)
at org.apache.hadoop.hbase.coprocessor.CoprocessorHost$Observer0
perationWithoutResult.callObserver(CoprocessorHost.java:558)
at org.apache.hadoop.hbase.coprocessor.CoprocessorHost.execOpe
ration(CoprocessorHost.java:631)
at org.apache.hadoop.hbase.master.MasterCoprocessorHost.preRev
oke(MasterCoprocessorHost.java:1854)
at org.apache.hadoop.hbase.master.MasterRpcServices.revoke(Mas
terRpcServices.java:2740)
at org.apache.hadoop.hbase.shaded.protobuf.generated.MasterPr
otos$MasterService$2.callBlockingMethod(MasterProtos.java)
at org.apache.hadoop.hbase.ipc.RpcServer.call(RpcServer.java:387
)
at org.apache.hadoop.hbase.ipc.CallRunner.run(CallRunner.java
:139)
at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExe
cutor.java:369)
at org.apache.hadoop.hbase.ipc.RpcExecutor$Handler.run(RpcExe
cutor.java:349)
```

None

**CDPD-56741: Improvement in log message when jwtauth not used**

The following exception is printed at startup only and it is not cluttering logs:

```
2023-05-30 06:18:40,127 ERROR org.apache.ranger.rms.security.RMS
JwtAuthFilter:
quasar-pibgzl-1.quasar-pibgzl.root.hwx.site-startStop-1]:
Failed to initialize Ranger RMS JWT Auth Filter.
java.lang.Exception: RangerJwtAuthHandler:
Mandatory configs ('jwks.provider-url' & 'jwt.public-key') are m
issing, must provide atleast one.
at org.apache.ranger.authz.handler.jwt.RangerJwtAuthHandler.in
italize(RangerJwtAuthHandler.java:84)
~[ranger-authn-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
at org.apache.ranger.rms.security.RMSJwtAuthFilter.initialize(R
MSJwtAuthFilter.java:77)
~[ranger-rms-common-2.4.0.7.1.9.0-186.jar:2.4.0.7.1.9.0-186]
at jdk.internal.reflect.NativeMethodAccessorImpl.invoke0(Native
Method) ~[?:?]
at jdk.internal.reflect.NativeMethodAccessorImpl.invoke(NativeMe
thodAccessorImp
```

This exception is logged, since mandatory JWT auth filter configs are not provided (expected for some environments like y-cloud) Even though JWT auth filter is failed to initialised, it falls back to kerberos auth, hence no impact from auth perspective

NA



**CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup**

This is a warning log printed in catalina.out file when Ranger RMS server is initialised. Following exception is observed only in Oracle 19 setup. FileNotFoundException: /usr/share/java/oraclepki.jar  
NA

**CDPD-55107: Not able to search using multiple user filter in access audit tab**

If you were using multiple user search filters in Audit > Access Tab on Ranger Admin UI, after upgrading to CDP-7.1.9 that would not be supported. You can continue to search user with single search filter.

None

**CDPD-48975: Ranger KMS KTS to KMS DB migration : keys with the same name but different case are not migrated**

KMS keys are not case sensitive

No work arounds. Such keys combination are very rare and migration doc was updated to check such keys before starting the migration.

**CDPD-58704: hadoop roll key / key delete command shows operation failed error when one KMS host is down, even when operation succeeds**

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered kms instances. if one KMS instance is stopped ( not deleted), client gets runtime exception.

This simply returns the runtime exception on client end for stopped instances but doesn't break any functionality.

**CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload:**

```
{"resourceName": "classification", "userInput": "", "resources": { "classification": [] }}
```

expectation is to return all the classifications . But the response has only "classification" Happens similarly for entity-label , entity-business-metadata.

None.

**CDPD-42598: Kafka policy creation allowed with incorrect permissions.**

When creating a Kafka policy from the UI, the permissions "Idempotent write" and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource

None.

**CDPD-40734: User allowed to insert data into a hive table when there is a deny policy on a table column.**

A user is allowed to enter data into a table even if there is a deny policy present on one of the table columns.

Test scenario details:

```
Policy setup :-
policy 1 :- all access policy for hrt_qa, hive and impala users
resources - database - * , table - * , column - *
users : hrt_qa, hive, impala
access - all access allowed
policy 2 :- policy on test_1.table_1 for hrt_5
users : hrt_5
resources : database - test_1, table - table_1, column - *
```

```

access :- all access allowed
policy 3 :- deny policy on test_1.table_1.c0 for hrt_5
users : hrt_5
resources : database - test_1, table - table_1, column - c0
access - all access denied
data setup :-
database - test_1
table - table_1(c0 int, c1 int)

```

The user is able to insert data into the table.

None.

**CDPD-58860: After upgrading CDP-7.1.8 to CDP-7.1.9, cdp-proxy token missing from Knox Ranger policy.**

As part of [OPSAPS-67480](#) in 719 default ranger policy is added from cdp-proxy-token topology, so that after a new installation of CDP-7.1.9, the Knox Ranger policy includes cdp-proxy-token. However, upgrades do not add cdp-proxy-token to cm\_knox policies automatically.

Manually add cdp-proxy-token to the Knox policy, using Ranger Admin Web UI.

1. Log in to Cloudera Manager Ranger Admin Web UI, as a Ranger administrator.
2. On Ranger Admin Web UI Service Manager Resource Knox, click cm\_knox.
3. In Knox Policies, open the CDP Proxy UI, API and Token policy.
4. In Knox Topology\*, add cdp-proxy-token.
5. Click Save.
6. Restart Ranger.

**CDPD-60633: User with user-role is revoked default 'Security Zone' module permission from Ranger Admin, will not be able to access policies on Ranger Admin UI**

By default, user with user role has permission to 'Security Zone' module. If you have revoked the permission of a user with user role for 'Security Zone' module from Ranger Admin, then the policy listing page is stuck on loading.

As a Ranger Administrator user, edit the User and Group permissions for Security Zone by adding the user with user-role to the Security Zone, as follows:

1. Log in to Cloudera Manager Ranger Admin Web UI, as a Ranger administrator.
2. On Ranger Admin Web UI Settings Permissions Security Zone, click Edit.
3. On Edit Permission User and Group Permissions Select and Add User :
  - a. Click Select Users.
  - b. Choose {User}.
  - c. Click +.
  - d. Click Save.
  - e. Log out and log in again with the added user.

If you have a restriction to not give the permission to the user, then there is an alternate workaround to switch to Ranger Admin using Backbone JS where the user with restricted 'Security Zone' module permission can access policies.

1. Login to Ranger Admin Web UI.
2. On Ranger Admin Web UI User Profile menu, click Backbone Classic UI.

You will be redirected to the old Ranger Admin UI where you can access policies.

**CDPD-61439: In Tag-based policy from Ranger Admin UI, Allow Conditions permissions item is not showing services permissions which have enableDenyAndExceptionsInPolicies flag false**

Steps to reproduce :

1. On Ranger Admin Web UI Service Manager Tag Policies Access tab, click Add New Policy.

2. On Create Policy Allow Conditions (component) Permissions , click +.

In Component Permissions Select Component ; the following components, which have the option `enableDenyAndExceptionsInPolicies=false`, do not appear listed:

- elasticsearch
- kylin
- nifi-registry
- nifi
- sqoop

However, these service components should be shown in the Allow condition.

Use the classic UI, as follows:

1. Log in to Cloudera Manager Ranger Admin Web UI , as a Ranger administrator.
2. On Ranger Admin Web UI User Profile menu , click Backbone Classic UI.
3. Go to Access Manager Tag Based Policies Access Add New Policy .
4. In Allow Conditions Component Permissions , click +.

All components for which `enableDenyAndExceptionsInPolicies=false`, appear listed.

## Known Issues in Ranger KMS

Learn about the known issues in Ranger KMS, the impact or changes to the functionality, and the workaround.

### CDPD-70115: Ranger KMS with Oracle DB not supported for Navigator Encrypt

Navigator Encrypt deposit registration is failing with Ranger KMS DB with Oracle DB setup with the following error:

```
java.sql.SQLException: ORA-02289: sequence does not exist
Error Code: 2289
```

None.

## Known Issues in Schema Registry

Learn about the known issues in Schema Registry, the impact or changes to the functionality, and the workaround.

### CDPD-40380: Authorization checking issue when Kerberos is disabled

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the `ranger.plugin.schema-registry.service.name` field.
3. Replace `GENERATED_RANGER_SERVICE_NAME` with the actual name of the service.
4. Restart the Schema Registry service.

### CDPD-49304: AvroConverter does not support composite default values

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

### CDPD-60160: Schema Registry Atlas integration does not work with Oracle databases

Schema Registry is unable to create entities in Atlas if Schema Registry uses an Oracle database. The following will be present in the Schema Registry log if you are affected by this issue:

```
ERROR com.cloudera.dim.atlas.events.AtlasEventsProcessor: An error occurred while processing Atlas events.
java.lang.IllegalArgumentException: Cannot invoke com.hortonworks.registries.schemaregistry.AtlasEventStorable.setType on bean class 'class com.hortonworks.registries.schemaregistry.AtlasEventS
```

```
torable' - argument type mismatch - had objects of type "java.lang.Long" but expected signature "java.lang.Integer"
```

This issue causes the loss of audit data on Oracle environments.

None.

#### **CDPD-59015: Schema Registry does not create new versions of schemas even if the schema is changed**

Schema Registry uses a schema fingerprinting mechanism to differentiate between schemas. However, fingerprinting does not take into consideration the schema attributes of the field type. As a result, if you have two schemas where the only difference is that one has type attributes defined and the other does not, they will be considered identical by Schema Registry. For example, the following schemas are considered identical:

```
#Schema V1
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [
 {"name": "local_timestamp_micros_long", "type": "long"}]}

#Schema V2
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [
 {"name": "local_timestamp_micros_long", "type": {"type": "long", "logicalType": "local-timestamp-micros"}}]}
```

Notice that the only difference is that in the second schema, the `local_timestamp_micros_long` field has a logical type specified. In cases like this, the new version of the schema is not created, the initial version is used. This is true even if the data that is being produced has a new schema version. The ID of the first schema version is used and is put in the serialized record. The new schema version is not created.

This issue is common when using change data capture (CDC) connectors like the Debezium connectors. This is because CDC connectors create schemas with the logical type decimal based on the column type in the database schema. For example:

```
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [
 {"name": "database_column", "type": {"type": "bytes", "logicalType": "decimal", "precision": 64, "scale": 0}}]}
```

If the database schema changes (for example, the column type), it is possible that only scale changes, which is a schema attribute.

```
{"type": "record", "name": "schema_name", "namespace": "ns", "fields": [
 {"name": "database_column", "type": {"type": "bytes", "logicalType": "decimal", "precision": 64, "scale": 1}}]}
```

In this case, even though scale changed to 1, the first version of the schema is used where scale is 0. As a result, the data is consumed with the wrong scale.

Avoid using logical types or other attributes. Alternatively, ensure that there are no changes in the logical types or other attributes between schema versions.

#### **OPSAPS-68708: Schema Registry might fail to start if a load balancer address is specified in Ranger**

Schema Registry does not start if the address specified in the Load Balancer Address Ranger property does not end with a trailing slash (/).

Set the value of the `RANGER_REST_URL` Schema Registry environment variable to an address that includes a trailing slash.

1. In Cloudera Manager, select the Schema Registry service.
2. Go to Configuration.

- Find the Schema Registry Server Environment Advanced Configuration Snippet (Safety Valve) property and add the following:

```
Key: RANGER_REST_URL
Value: [***RANGER REST API URL***]
```

Replace [\*\*\*RANGER REST API URL\*\*\*] with an address that can be used by Schema Registry to access Ranger. Ensure that the address ends with a trailing slash. For example: `http://ranger-1.cloudera.com:6182/`

- Restart the Schema Registry service.

#### **OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade**

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

- Access the Ranger Console (Ranger Admin web UI).
- Click the `cm_atlas` resource-based service.
- Add the `schemaregistry` user to the all - \* policies.
- Click `Manage Service Edit Service`.
- Add the `schemaregistry` user to the `default.policy.users` property.

### **Known Issues in Search Client**

Learn about the known issues in Search Client, the impact or changes to the functionality, and the workaround.

#### **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.

### **Known Issues in Apache Solr**

Learn about the known issues in Apache Solr, the impact or changes to the functionality, and the workaround.

#### **HBase Indexer does not work with JDK 17**

Depending on the Cloudera Manager version used with CDP, HBase Indexer (KS Indexer) may have compatibility issues with JDK 17.

You have the following options to fix this issue:

- Upgrade Cloudera Manager to version 7.11.3 or higher.
- If upgrading Cloudera Manager is not an option, you can manually add the following to HBase Indexer Java options in Cloudera Manager:

```
--add-opens java.base/java.nio=ALL-UNNAMED --add-opens java.
base/java.util.concurrent.atomic=ALL-UNNAMED --add-opens jav
a.base/java.lang=ALL-UNNAMED --add-opens java.base/java.lang
.reflect=ALL-UNNAMED
```

#### **Splitshard operation fails after CDH 6 to CDP upgrade**

Collections are not reindexed during an upgrade from CDH 6 to CDP 7 because Lucene 8 (CDP) can read Lucene 7 (CDH 6) indexes.

If you try to execute a SPLITSHARD operation against such a collection, it fails with a similar error message:

```
o.a.s.h.a.SplitOp ERROR executing split: => java.lang.IllegalArg
umentException: Cannot merge a segment t
hat has been created with major version 7 into this index which
has been created by major version 8
 at org.apache.lucene.index.IndexWriter.validateMergeRea
der(IndexWriter.java:3044)
```

```

java.lang.IllegalArgumentException: Cannot merge a segment that has
been created with major version 7 into this index which has been
created by major version 8
 at org.apache.lucene.index.IndexWriter.validateMergeReader(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(IndexWriter.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(SolrIndexSplitter.java:318) ~[solr-core-8.11.2.7.1.9.3-2.jar:8.11.2.7.1.9.3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins - 2023-12-02 00:16:28]
 at org.apache.solr.update.SolrIndexSplitter.split(SolrIndexSplitter.java:184) ~[solr-core-8.11.2.7.1.9.3-2.jar:8.11.2.7.1.9.3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins - 2023-12-02 00:16:28]
 at org.apache.solr.update.DirectUpdateHandler2.split(DirectUpdateHandler2.java:922) ~[solr-core-8.11.2.7.1.9.3-2.jar:8.11.2.7.1.9.3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins - 2023-12-02 00:16:28]
 at org.apache.solr.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 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins - 2023-12-02 00:16:28]
 at org.apache.solr.handler.admin.CoreAdminOperation.execute(CoreAdminOperation.java:367) ~[solr-core-8.11.2.7.1.9.3-2.jar:8.11.2.7.1.9.3-2 a6ff93f9665115dffbdad0ad7f222fd1978d495d - jenkins - 2023-12-02 00:16:28]

```

This happens because the segment created using a Lucene 7 index cannot be merged into a Lucene 8 index.

Drop the entire collection, delete the data in HDFS and recreate the collection with Solr 8 configs.

### Changing the default value of Client Connection Registry HBase configuration parameter causes HBase MRIT job to fail

If the value of the HBase configuration property Client Connection Registry is changed from the default ZooKeeper Quorum to Master Registry then the Yarn job started by HBase MRIT fails with a similar error message:

```

Caused by: org.apache.hadoop.hbase.exceptions.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'
```

### Solr does not support rolling upgrade to release 7.1.9 or lower

Solr supports rolling upgrades from release 7.1.9 and higher. Upgrading from a lower version means that all the Solr Server instances are shut down, parcels upgraded and activated and then the Solr Servers are started again. This causes a service interruption of several minutes, the actual value depending on cluster size.

Services like Atlas and Ranger that depend on Solr, may face issues because of this service interruption.

None.

### Unable to see single valued and multivalued empty string values when querying collections after upgrade to CDP

After upgrading from CDH or HDP to CDP, you are not able to see single valued and multi Valued empty string values in CDP.

This behavior in CDP is due to the remove-blank processor present in solrconfig.xml in Solr 8.

Remove the remove-blank processor from solrconfig.xml.

### Cannot create multiple heap dump files because of file name error

Heap dump generation fails with a similar error message:

```

java.lang.OutOfMemoryError: Java heap space
Dumping heap to /data/tmp/solr_solr-SOLR_SERVER-fc9dacc265fabfc500b92112712505e3_pid{{PID}}.hprof ...
Unable to create /data/tmp/solr_solr-SOLR_SERVER-fc9dacc265fabfc500b92112712505e3_pid{{PID}}.hprof: File exists

```

The cause of the problem is that {{PID}} does not get substituted during dump file creation with an actual process ID and because of that, a generic file name is generated. This causes the next dump file creation to fail, as the existing file with the same name cannot be overwritten.

You need to manually delete the existing dump file.

### **Solr coreAdmin status throws Null Pointer Exception**

You get a Null Pointer Exception with a similar stacktrace:

```
Caused by: java.lang.NullPointerException
 at org.apache.solr.core.SolrCore.getInstancePath(SolrCore.java:333)
 at org.apache.solr.handler.admin.CoreAdminOperation.getCoreStatus(CoreAdminOperation.java:324)
 at org.apache.solr.handler.admin.StatusOp.execute(StatusOp.java:46)
 at org.apache.solr.handler.admin.CoreAdminOperation.execute(CoreAdminOperation.java:362)
```

This is caused by an error in handling solr admin core STATUS after collections are rebuilt.

Restart the Solr server.

### **Applications fail because of mixed authentication methods within dependency chain of services**

Using different types of authentication methods within a dependency chain, for example, configuring your indexer tool to authenticate using Kerberos and configuring your Solr Server to use LDAP for authentication may cause your application to time out and eventually fail.

Make sure that all services in a dependency chain use the same type of authentication.

### **API calls fail with error when used with alias, but work with collection name**

API calls fail with a similar error message when used with an alias, but they work when made using the collection name:

```
[] o.a.h.s.t.d.w.DelegationTokenAuthenticationFilter Authentication exception: User: xyz@something.example.com is not allowed to impersonate xyz@something.example.com
[c:RTOTagMetaOdd s:shard3 r:core_node11 x:RTOTagMetaOdd_shard3_replica_n8] o.a.h.s.t.d.w.DelegationTokenAuthenticationFilter Authentication exception: User: xyz@something.example.com is not allowed to impersonate xyz@something.example.com
```

Make sure there is a replica of the collection on every host.

### **CrunchIndexerTool does not work out of the box if /tmp is mounted noexec mode**

When you try to run CrunchIndexerTool with the /tmp directory mounted in noexec mode, It throws a snappy-related error.

Create a separate directory for snappy temp files which is mounted with EXEC privileges and set this directory as the value of the org.xerial.snappy.tmpdir java property as a driver java option.

For example:

```
export myDriverJarDir=/opt/cloudera/parcels/CDH//lib/solr/contrib/crunch;export myDependencyJarDir=/opt/cloudera/parcels/CDH//lib/search/lib/search-crunch;export myDriverJar=$(find $myDriverJarDir -maxdepth 1 -name 'search-crunch-*.jar' ! -name '*-job.jar' ! -name '*-sources.jar');export myDependencyJarFiles=$(find $myDependencyJarDir -name '*.jar' | sort | tr '\n' ',' | head -c -1);export myDependencyJarPaths=$(find $myDependencyJarDir -name '*.jar' | sort | tr '\n' ':' | head -c -1);export HADOOP_CONF_DIR=;spark-submit --master local --deploy-mode client --driver-library-path /opt/cloudera/parcels/CDH//lib/hadoop/lib/
```



```
native/ --jars $myDependencyJarFiles --driver-java-options ' -
Dorg.xerial.snappy.tmpdir=/home/systest/tmp ' --class org.apa
che.solr.crunch.CrunchIndexerTool $myDriverJar --input-file-form
at=avroParquet --input-file-reader-schema search-parquetfile/par
quet-schema.avsc --morphline-file /tmp/mrTestBase.conf --pipelin
e-type spark --chatty hdfs://[***HOSTNAME***]:8020/tmp/parquetfi
leparsertest-input
```

### Mergeindex operation with --go-live fails after CDH 6 to CDP upgrade

During an upgrade from CDH6 to CDP, collections are not reindexed because Lucene 8 (CDP) can read Lucene 7 (CDH6) indexes.

If you try to execute MapReduceIndexerTool (MRIT) or HBase Indexer MRIT with --go-live against such a collection, you get a similar error message:

```
Caused by: java.lang.IllegalArgumentException: Cannot merge a se
gment 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.validateMergeReade
r(IndexWriter.java:2894)
 at org.apache.lucene.index.IndexWriter.addIndexes(Index
Writer.java:2960)
 at org.apache.solr.update.DirectUpdateHandler2.mergeIn
dexes(DirectUpdateHandler2.java:570)
 at org.apache.solr.update.processor.RunUpdateProcessor.
processMergeIndexes(RunUpdateProcessorFactory.java:95)
 at org.apache.solr.update.processor.UpdateRequestProcesso
r.processMergeIndexes(UpdateRequestProcessor.java:63)
```

This happens because CDP MRIT and HBase indexer use Solr 8 as embedded Solr, which creates a Lucene 8 index. It cannot be merged (using MERGEINDEXES) into an older Lucene 7 index.

In the case of MRIT the only way to move past this issue is to drop the entire collection, delete the data in HDFS and recreate the collection with Solr 8 configs.

For HBase Indexer MRIT an alternative workaround is setting the number of reducers to 0 (--reducers 0) because in this case documents are sent directly from the mapper tasks to live Solr servers instead of using MERGEINDEXES.

### Apache Tika upgrade may break morphlines indexing

The upgrade of Apache Tika from 1.27 to 2.3.0 brought potentially breaking changes for morphlines indexing. Duplicate/triplicate keys names were removed and certain parser class names were changed (For example, org.apache.tika.parser.jpeg.JpegParser changed to org.apache.tika.parser.image.JpegParser).

To avoid morphline commands failing after the upgrade, do the following:

- Check if key name changes affect your morphlines. For more information, see *Removed duplicate/triplicate keys* in [Migrating to Tika 2.0.0](#).
- Check if the name of any parser you use has changed. For more information, see the Apache Tika [API documentation](#).

Update your morphlines if necessary.

### CDPD-28006: Solr access via Knox fails with impersonation error though auth\_to\_local and proxy user configs are set

Currently the names of system users which are impersonating users with Solr should match with the names of their respective Kerberos principals.

If, for some reason, this is not feasible, you must add the user name you want to associate with the custom Kerberos principal to Solr configuration via the Solr Service Environment Advanced Configuration Snippet (Safety Valve) environment variable in Cloudera Manager.

For more information, see [Configuring custom Kerberos principals and custom system users](#).

#### **CDH-77598: Indexing fails with socketTimeout**

Starting from CDH 6.0, the HTTP client library used by Solr has a default socket timeout of 10 minutes. Because of this, if a single request sent from an indexer executor to Solr takes more than 10 minutes to be serviced, the indexing process fails with a timeout error.

This timeout has been raised to 24 hours. Nevertheless, there still may be use cases where even this extended timeout period proves insufficient.

If your MapreduceIndexerTool or HBaseMapreduceIndexerTool batch indexing jobs fail with a timeout error during the go-live (Live merge, MERGEINDEXES) phase (This means the merge takes longer than 24 hours).

Use the `--go-live-timeout` option where the timeout can be specified in milliseconds.

#### **CDPD-12450: CrunchIndexerTool Indexing fails with socketTimeout**

The http client library uses a socket timeout of 10 minutes. The Spark Crunch Indexer does not override this value, and in case a single batch takes more than 10 minutes, the entire indexing job fails. This can happen especially if the morphlines contain DeleteByQuery requests.

Try the following workarounds:

- Check the batch size of your indexing job. Sending too large batches to Solr might increase the time needed on the Solr server to process the incoming batch.
- If your indexing job uses deleteByQuery requests, consider using deleteById wherever possible as deleteByQuery involves a complex locking mechanism on the Solr side which makes processing the requests slower.
- Check the number of executors for your Spark Crunch Indexer job. Too many executors can overload the Solr service. You can configure the number of executors by using the `--mappers` parameter
- Check that your Solr installation is correctly sized to accommodate the indexing load, making sure that the number of Solr servers and the number of shards in your target collection are adequate.
- The socket timeout for the connection can be configured in the morphline file. Add the `solrClientSocketTimeout` parameter to the `solrLocator` command

Example

```
SOLR_LOCATOR :
{
 collection : test_collection
 zkHost : "zookeeper1.example.corp:2181/solr"
 # 10 minutes in milliseconds
 solrClientSocketTimeout: 600000
 # Max number of documents to pass per RPC from morphline to
 Solr Server
 # batchSize : 10000
}
```

#### **CDPD-29289: HBaseMapReduceIndexerTool fails with socketTimeout**

The http client library uses a socket timeout of 10 minutes. The HBase Indexer does not override this value, and in case a single batch takes more than 10 minutes, the entire indexing job fails.

You can overwrite the default 600000 millisecond (10 minute) socket timeout in HBase indexer using the `--solr-client-socket-timeout` optional argument for the direct writing mode (when the value of the `--reducers` optional argument is set to 0 and mappers directly send the data to the live Solr).

#### **CDPD-20577: Splitshard operation on HDFS index checks local filesystem and fails**

When performing a shard split on an index that is stored on HDFS, SplitShardCmd still evaluates free disk space on the local file system of the server where Solr is installed. This may cause the command to fail, perceiving that there is no adequate disk space to perform the shard split.

Run the following command to skip the check for sufficient disk space altogether:

- On nonsecure clusters:

```
curl 'http://$[***SOLR_SERVER_HOSTNAME***]:8983/solr/admin/collections?action=SPLITSHARD&collection=[***COLLECTION_NAME***]&shard=[***SHARD_TO_SPLIT***]&skipFreeSpaceCheck=true'
```

- On secure clusters:

```
curl -k -u : --negotiate 'http://$[***SOLR_SERVER_HOSTNAME***]:8985/solr/admin/collections?action=SPLITSHARD&collection=[***COLLECTION_NAME***]&shard=[***SHARD_TO_SPLIT***]&skipFreeSpaceCheck=true'
```

Replace `[***SOLR_SERVER_HOSTNAME***]` with a valid Solr server hostname, `[***COLLECTION_NAME***]` with the collection name, and `[***SHARD_TO_SPLIT***]` with the ID of the to split.

To verify that the command executed successfully, check overseer logs for a similar entry:

```
2021-02-02 12:43:23.743 INFO (OverseerThreadFactory-9-thread-5-processing-n:myhost.example.com:8983_solr) [c:example s:shard1] o.a.s.c.a.c.SplitShardCmd Skipping check for sufficient disk space
```

#### DOCS-5717: 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.

None, you need to reindex collections.

#### 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.

#### CDH-19407: The Browse and Spell Request Handlers are not enabled in schemaless mode

The Browse and Spell Request Handlers require certain fields to be present in the schema. Since those fields cannot be guaranteed to exist in a Schemaless setup, the Browse and Spell Request Handlers are not enabled by default.

If you require the Browse and Spell Request Handlers, add them to the solrconfig.xml configuration file. Generate a non-schemaless configuration to see the usual settings and modify the required fields to fit your schema.

#### **CDH-17978: 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 indexes. Blockcache writing is disabled by default.

None

#### **CDH-58276: Users with insufficient Solr permissions may receive a "Page Loading" message from the 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 receive a web page that never finishes loading.

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.

### **Unsupported features**

The following Solr features are currently not supported in Cloudera Data Platform:

- Panel with security info in admin UI's dashboard
- Incremental backup mode
- Schema Designer UI
- Package Management System
- HTTP/2

- Solr SQL/JDBC
- Graph Traversal
- Cross Data Center Replication (CDCR)
- SolrCloud Autoscaling
- HDFS Federation
- Saving search results
- Solr contrib modules (Spark, MapReduce and Lily HBase indexers are not contrib modules but part of the Cloudera Search product itself, therefore they are supported).

### Limitations

#### 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.

### Known Issues in Apache Spark

Learn about the known issues in Spark, the impact or changes to the functionality, and the workaround.

#### CDPD-60862: Rolling restart fails during ZDU when DDL operations are in progress

During a Zero Downtime Upgrade (ZDU), the rolling restart of services that support Data Definition Language (DDL) statements might fail if DDL operations are in progress during the upgrade. As a result, ensure that you do not run DDL statements during ZDU.

The following services support DDL statements:

- Impala
- Hive – using HiveQL
- Spark – using SparkSQL
- HBase
- Phoenix
- Kafka

Data Manipulation Language (DML) statements are not impacted and can be used during ZDU. Following the successful upgrade, you can resume running DDL statements.

None. Cloudera recommends modifying applications to not use DDL statements for the duration of the upgrade. If the upgrade is already in progress, and you have experienced a service failure, you can remove the DDLs in-flight and resume the upgrade from the point of failure.

#### CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the RuntimeWarning: Failed to add file message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

#### CDPD-67517: Spark3 tests are failing if /tmp is mounted as NOEXEC.

Map the tmpdir to a writable path in spark3-conf/spark-defaults.conf using the following steps:

1. In the Cloudera Data Platform (CDP) Management Console, go to Data Hub Clusters.
2. Find and select the cluster you want to configure.
3. Click the link for the Cloudera Manager URL.
4. Go to the Spark service.
5. Click the Configuration tab.

6. Select Scope Gateway .
7. Select Category Advanced .
8. Locate the Spark Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-defaults.conf\_client\_config\_safety\_valve property.
9. Map the tmpdir to a writable path:

```
spark.driver.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
spark.executor.extraJavaOptions=-Djava.io.tmpdir=/var/tmp
```

10. Enter a Reason for change, and then click Save Changes to commit the changes.
11. Deploy the client configuration.

**CDPD-23817: In the upgraded Cluster, the permission of /tmp/spark is restricted due to the HDP configuration hive.exec.scratchdir=/tmp/spark.**

If you are using the /tmp/spark directory in the CDP cluster, you must provide the required additional Policies/ACL permissions.

**CDPD-22670 and CDPD-23103: There are two configurations in Spark, "Atlas dependency" and "spark\_lineage\_enabled", which are conflicted. The issue is when Atlas dependency is turned off but spark\_lineage\_enabled is turned on.**

Run Spark application, Spark will log some error message and cannot continue. That can be restored by correcting the configurations and restarting Spark component with distributing client configurations.

**CDPD-23007: Mismatch in the Spark Default DB Location. In HDP 3.1.5, hive\_db entities have one attribute - 'location' which is configured to the '/managed' path. In fresh install of CDP 7.1.7, hive\_db entities now have 2 attributes 'location' configured to '/external' path and 'managedLocation' configured to '/managed' path. In the AM2CM migration (HDP 3.1.5 -> CDP 7.1.7), the 'location' attribute from hive\_db entities in HDP 3.1.5 comes unaltered to CDP 7.1.7 and hence maps to '/managed' path.**

This issue arises only if you are upgrading from HDP 3.1.5 to CDP 7.1.7. If you are performing a fresh install of CDP 7.1.7, you can ignore this issue.

None

**CDPD-217: The Apache Spark connector is not supported**

The old *Apache Spark - Apache HBase Connector* (shc) is not supported in CDP releases.

Use the new HBase-Spark connector shipped in CDP release.

**CDPD-3038: Launching pyspark displays several HiveConf warning messages**

When pyspark starts, several Hive configuration warning messages are displayed, similar to the following:

```
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of name hive.meta
store.runworker.in does not exist
23/08/02 08:37:26 WARN conf.HiveConf: HiveConf of name hive.ma
sking.algo does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of name hive.me
tastore.runworker.in does not exist
23/08/02 08:37:34 WARN conf.HiveConf: HiveConf of name hive.mask
ing.algo does not exist
```

These errors can be safely ignored.

## 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****CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic**

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are a 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

**CDPD-11079: Blacklisted topics appear in the list of replicated topics**

If a topic was originally replicated but was later excluded for replication, 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, this topic will be included in the response. Additionally, the excluded topic will also be visible in the SMM UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

**CDPD-30275: SRM may automatically re-create deleted topics on target clusters**

If `auto.create.topics.enable` is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

1. Remove the topic from the topic allowlist with `srm-control`. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGE
T_CLUSTER] --remove
[TOPIC1]
```

2. Wait until SRM is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

**OPSAPS-71258: Zstandard and Snappy compression do not support /tmp mounted as noexec**

SRM cannot process messages compressed with Zstandard or Snappy if /tmp is mounted as noexec.

The workaround steps for Zstandard compression are the following. You need to complete these steps for all the collected hosts.

1. Identify the hosts running SRM.
2. Verify the Zstandard version by checking the classpath of the running process.

```
ps aux
```

By default, SRM uses version `libzstd-jni-1.5.2-1`.

The Zstandard JAR file is located at `/opt/cloudera/parcels/CDH/jars/zstd-jni-1.5.2-1.jar`. This file includes the architecture and OS-specific binaries. Check [Zstandard JNI GitHub repository](#) for the full list of supported native libraries.

3. Extract the binary corresponding to the target system similar to the following example.

```
/usr/java/default/bin/jar xf /opt/cloudera/parcels/CDH/jars/
zstd-jni-1.5.2-1.jar linux/amd64/libzstd-jni-1.5.2-1.so
```

4. Copy the extracted binary to a location included in `java.library.path` similar to the following example.

```
cp linux/amd64/libzstd-jni-1.5.2-1.so /lib
```

5. Restart SRM.

The workaround steps for Snappy compression are the following:

1. In Cloudera Manager, select the SRM service and go to Configuration.
2. Add the following line to SRM\_JVM\_PERF\_OPTS.

```
-org.xerial.snappy.tmpdir==[***PATH***]
```

Where `[***PATH***]` is a directory that is not `/tmp`.

3. Restart SRM.

#### **CDPD-77283: SRM can skip records when replication-records-lag endoffset fetching fails**

SRM can skip records during data replication when end offset fetching fails.

None.

#### **CDPD-80872: SRM replication-records-lag is incorrect on empty partitions with non-zero end offset**

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**

#### **SRM cannot replicate Ranger authorization policies to or from Kafka clusters**

Due to a limitation in the Kafka-Ranger plugin, SRM cannot replicate Ranger policies to or from clusters that are configured to use Ranger for authorization. If you are using SRM to replicate data to or from a cluster that uses Ranger, disable authorization policy synchronization in SRM. This can be achieved by clearing the Sync Topic Acls Enabled (`sync.topic.acls.enabled`) checkbox.

#### **SRM cannot ensure the exactly-once semantics of transactional source topics**

SRM data replication uses at-least-once guarantees, and as a result cannot ensure the exactly-once semantics (EOS) of transactional topics in the backup/target cluster.



**Note:** Even though EOS is not guaranteed, you can still replicate the data of a transactional source, but you must set `isolation.level` to `read_committed` for SRM's internal consumers. This can be done by adding `[***SOURCE CLUSTER ALIAS***]->[***TARGET CLUSTER ALIAS***].consumer.isolation.level=read_committed` to the Streams Replication Manager's Replication Configs SRM service property in Cloudera Manager. The `isolation.level` property can be set on a global connector or replication level. For example:

```
#Global connector level
connectors.consumer.isolation.level=read_committed
#Replication level
uswest->useast.consumer.isolation.level=read_committed
```

#### **SRM checkpointing is not supported for transactional source topics**

SRM does not correctly translate checkpoints (committed consumer group offsets) for transactional topics. Checkpointing assumes that the offset mapping function is always increasing, but with transactional source topics this is violated. Transactional topics have control messages in them, which take up an offset in the log, but they are never returned on the consumer API. This causes the mappings to decrease, causing issues in the checkpointing feature. As a result of this limitation, failover operations for transactional topics is not possible.

### **Known Issues for Apache Sqoop**

Learn about the known issues in Sqoop, the impact or changes to the functionality, and the workaround.

#### **CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version**



If you have upgraded to CDP PvC Base 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to CDP PvC Base 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see [Troubleshooting Apache Sqoop issues](#).

#### **CDPD-44431: Using direct mode causes problems**

Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Stop using direct mode. Do not use the `--direct` option in Sqoop import or export commands.

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the `sqoop.enable.deprecated.direct` property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through `-Dsqoop.enable.deprecated.direct=true`.

#### **CDPD-3089: Avro, S3, and HCat do not work together properly**

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

#### **Parquet columns inadvertently renamed**

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.

#### **PARQUET-99: Importing Parquet files might cause out-of-memory (OOM) errors**

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.

## **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.

#### **CDPD-70906: NullPointerException when opening SMM Web UI home page without Cruise Control permissions**

When a user opens the SMM Web UI home page, a `NullPointerException` (NPE) error message is displayed in the top right corner of the screen. This occurs when the logged-in user does not have the necessary permissions to access Cruise Control. Instead of a clear authorization error, the application throws an unhandled exception.

Grant the user access to Cruise Control by adding them to one of the three authorization properties: `ADMIN Level Users`, `USER Level Users`, or `VIEWER Level Users`.

#### **OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager**

Cloudera Manager does not display a Log Files menu for SMM UI role (and SMM UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by SMM UI is not supported by Cloudera Manager.

View the SMM UI logs on the host.

#### **OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners**

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL\_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in SMM safety valve in the following path:

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Add the following value for bootstrap servers>Save Changes > Restart SMM.

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

#### CDPD-46728: SMM UI shows the consumerGroup instead of the instances on the Profile page's right hand side

On the ConsumerGroupDetail page, SMM UI shows the group instead of its instances on the right hand side table.

None.

#### CDPD-47836: The FROM OFFSET field of the offset slider in Data Explorer does not update on partition change

When changing the partition on the **Data Explorer** tab of a topic, the **FROM OFFSET** field of the offset slider is not updated to reflect the first offset of the newly selected partition. If the first offset of the newly selected partition differs from the previous one, the application throws an error. This issue only affects the **Data Explorer** tab, which is accessed from the topic details page of a topic.

This issue does not affect the **Data Explorer** modal window, which is accessed by clicking  next to the name of a topic.

After selecting a new partition, refresh the page. Alternatively, use the **Data Explorer** modal window.

#### OPSAPS-71258: Zstandard and Snappy compression do not support /tmp mounted as noexec

SMM cannot process messages compressed with Zstandard or Snappy in Data Explorer if /tmp is mounted as noexec.

The workaround steps for Zstandard compression are the following. You need to complete these steps for all the collected hosts.

1. Identify the hosts running SMM.
2. Verify the Zstandard version by checking the classpath of the running process.

```
ps aux
```

By default, SMM uses version libzstd-jni-1.5.2-1.

The Zstandard JAR file is located at /opt/cloudera/parcels/CDH/jars/zstd-jni-1.5.2-1.jar. This file includes the architecture and OS-specific binaries. Check [Zstandard JNI GitHub repository](#) for the full list of supported native libraries.

3. Extract the binary corresponding to the target system similar to the following example.

```
/usr/java/default/bin/jar xf /opt/cloudera/parcels/CDH/jars/zstd-jni-1.5.2-1.jar linux/amd64/libzstd-jni-1.5.2-1.so
```

4. Copy the extracted binary to a location included in java.library.path similar to the following example.

```
cp linux/amd64/libzstd-jni-1.5.2-1.so /lib
```

5. Restart SMM.

The workaround steps for Snappy compression are the following:

1. In Cloudera Manager, select the SMM service and go to Configuration.
2. Add the following line to SMM\_JVM\_PERF\_OPTS.

```
-org.xerial.snappy.tmpdir==[***PATH***]
```

Where [\*\*\*PATH\*\*\*] is a directory that is not /tmp.

3. Restart SMM.

#### CDPD-82560: Streams Messaging Manager known unprotected endpoint

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

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

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 MapReduce and YARN

Learn about the known issues in MapReduce and YARN, the impact or changes to the functionality, and the workaround.

### Known Issues

#### COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

#### COMPX-13177: QueueManager webapp requests fail with 'HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS\_EDH\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA'

Products: Cloudera Manager for CDP Private Cloud Base Cloudera Manager for CDP Public Cloud Context: Centos 7.8 and Redhat 7.8 operating systems, when FIPS support is enabled. Problem: When attempting to display the Yarn Queue Manager interface, Cloudera Manager displays an error: "HTTP ERROR 400 java.net.ConnectException: Unsupported ciphersuite TLS\_EDH\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA". Workarounds: Edit the file /etc/default/cloudera-scm-server. Around line 28, modify the line that starts with "#export CMF\_OVERRIDE\_TLS\_CIPHERS=..." Remove the comment mark "#" Remove all ciphers with "3DES" in the name. Save the file. Re-start the Cloudera Manager Server service.

**CDPD-56559: MapReduce jobs can intermittently fail during a rolling upgrade.**

During a rolling upgrade between versions 7.1.9 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 7.1.9 to a higher version.

None

**COMPX-12021 Queue Manager configurations on Scheduler Configuration page are not working**

When setting the following properties on the YARN Queue Manager UI, the properties are set in `capacity-scheduler.xml` which does not have any effect on YARN. The properties need to be set in `yarn-site.xml`, 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-11380 Queue Manager displays an error stating that it is unable to complete a request**

If Queue Manager incorrectly detects an inactive Resource Manager, it displays an error stating that it is unable to complete a request.

Refresh the browser and reload the YARN Queue Manager UI.

**COMPX-6214: When there are more than 600 queues in a cluster, potential timeouts occur due to performance reasons that are visible in the Configuration Service.**

The Cloudera Manager proxy timeout configuration field is added now. This issue is tracked in OPSAPS-60554. For this release, the timeout is increased from 20 seconds to 5 minutes. However, if this problem occurs, Cloudera recommends you to increase the proxy timeout value.

**COMPX-5817: Queue Manager UI will not be able to present a view of pre-upgrade queue structure. CM Store is not supported and therefore Yarn will not have any of the pre-upgrade queue structure preserved.**

When a Data Hub cluster is deleted, all saved configurations are also deleted. All YARN configurations are saved in CM Store and this is yet to be supported in Data Hub and Cloudera Manager. Hence, the YARN queue structure also will be lost when a Data Hub cluster is deleted or upgraded or restored.

**COMPX-6628: Unable to delete single leaf queue assigned to a partition.**

In the current implementation, you cannot delete a single leaf queue assigned to a partition.

For each non-default partition, perform the following for the single child leaf queue and its parent queues:

1. In **Cloudera Manager**, click Cluster > YARN.
2. Click the **Configuration** tab.
3. Search for ResourceManager. In the Filters pane, under Scope, select ResourceManager.
4. Add the following in Capacity Scheduler Configuration Advanced Configuration Snippet (Safety Valve):

```
Name: yarn.scheduler.capacity.<queuePath>.accessible-node-labels.<partition>.capacity
Value: 0
```

Set the value to 0 in Percentage mode, and 0w in Weight mode, and [memory=0,vcores=0] in Absolute mode.

```
Name: yarn.scheduler.capacity.<queuePath>.accessible-node-labels.<partition>.maximum-capacity
Value: 100
```

Set the value to 100 in Percentage and Weight mode and [memory=0,vcores=0] in Absolute mode.

5. Adjust the capacities of the siblings of the parent queue for the same partition.
6. Click Save Changes.
7. Restart the active **ResourceManager** service for the changes to apply.
8. In **Cloudera Manager**, click Cluster > YARN Queue Manager UI.
9. Delete the desired single child leaf queue.

**COMPX-5240: Restarting parent queue does not restart child queues in weight mode**

When a dynamic auto child creation enabled parent queue is stopped in weight mode, its static and dynamically created child queues are also stopped. However, when the dynamic auto child creation enabled parent queue is restarted, its child queues remain stopped. In addition, the dynamically created child queues cannot be restarted manually through the YARN Queue Manager UI either.

Delete the dynamic auto child creation enabled parent queue. This action also deletes all its child queues, both static and dynamically created child queues, including the stopped dynamic queues. Then recreate the parent queue, enable the dynamic auto child creation feature for it and add any required static child queues.

**COMPX-5589: Unable to add new queue to leaf queue with partition capacity in Weight/Absolute mode**

Scenario

1. User creates one or more partitions.
2. Assigns a partition to a parent with children
3. Switches to the partition to distribute the capacities

4. Creates 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: Val
idation failed for Add queue
operation. Error message: CapacityScheduler configuration vali
dation failed:java.io.IOException:
Failed to re-init queues : Parent queue 'root.test2' have childr
en queue used mixed of weight
mode, percentage and absolute mode, it is not allowed, please do
uble check, details:
{Queue=root.test2.test2childNew, label= uses weight mode}. {Que
ue=root.test2.test2childNew,
label=partition uses percentage mode}
```

To create new queues under leaf queues without hitting this error, perform the following:

1. Switch to Relative mode
2. Create the required queues
3. Create the required partitions
4. Assign partitions and set capacities
5. Switch back to Weight mode
1. Create the entire queue structure
2. Create the required partitions
3. Assign partition to queues
4. Set partition capacities

#### **COMPX-5264: Unable to switch to Weight mode on creating a managed parent queue in Relative mode**

In the current implementation, if there is an existing managed queue in Relative mode, then conversion to Weight mode is not 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**

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 as the `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 dropdown list and distribute capacity to the queues before switching allocation mode or creating placement rules.

#### **COMPX-4992: Unable to switch to absolute mode after deleting a partition using YARN Queue Manager**

If you delete a partition (node label) which has been associated with queues and those queues have capacities configured for that partition (node label), the CS.xml still contains the partition (node label) information. Hence, you cannot switch to absolute mode after deleting the partition (node label).

It is recommended not to delete a partition (node label) which has been associated with queues and those queues have capacities configured for that partition (node label).

#### **COMPX-1445: Queue Manager operations are failing when Queue Manager is installed separately from YARN**

If Queue Manager is not selected during YARN installation, Queue Manager operation 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-1451: Queue Manager does not support multiple ResourceManagers**

When YARN High Availability is enabled there are multiple ResourceManagers. Queue Manager receives multiple ResourceManager URLs for a High Availability cluster. It picks the active ResourceManager URL only when Queue Manager page is loaded. Queue Manager cannot handle it gracefully when the currently active ResourceManager goes down while the user is still using the Queue Manager UI.

Reload the Queue Manager page manually.

**COMPX-3329: Autorestart is not enabled for Queue Manager in Data Hub**

In a Data Hub cluster, Queue Manager is installed with autorestart disabled. Hence, if Queue Manager goes down, it will not restart automatically.

If Queue Manager goes down in a Data Hub cluster, you must go to the Cloudera Manager Dashboard and restart the Queue Manager service.

**Third party applications do not launch if MapReduce framework path is not included in the client configuration**

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the `mapreduce.application.framework.path` property is set to the appropriate value, but third party applications with their own configurations will not launch.

Set the `mapreduce.application.framework.path` property to the appropriate configuration for third party applications.

**JobHistory URL mismatch after server relocation**

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated `mapred-site.xml` that references the correct JobHistory Server.

**CDH-49165: History link in ResourceManager web UI broken for killed Spark applications**

When a Spark application is killed, the history link in the ResourceManager web UI does not work.

To view the history for a killed Spark application, see the Spark HistoryServer web UI instead.

**CDH-6808: Routable IP address required by ResourceManager**

ResourceManager requires routable host:port addresses for `yarn.resourcemanager.scheduler.address`, and does not support using the wildcard `0.0.0.0` address.

Set the address, in the form `host:port`, either in the client-side configuration, or on the command line when you submit the job.

**YARN cannot start if Kerberos principal name is changed**

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN will not be able to start. In such case the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the `znode` and restart the YARN service.

- Use the reset ZK ACLs command. This also sets the znodes below /rmstore/ZKRMStateRoot to world:anyone:cdw which is less secure.

### **Queue Manager does not open on using a custom user with a default Kerberos principal**

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager process\_username property matches the YARN process\_username property.

### **COMPX-8687: Missing access check for getAppAttempts**

When the Job ACL feature is enabled using Cloudera Manager (YARN Configuration Enablg 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 will use the default value of this property. The default property of mapreduce.cluster.acls.enabled is false.

Workaround: Enable the Job ACL feature using an 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.

### **COMPX-7493: YARN Tracking URL that is shown in the command line does not work when Knox is enabled**

When Knox is configured for YARN, the Tracking URL printed in the command line of an YARN application such as spark-submit shows the direct URL instead of the Knox Gateway URL.

Upgrade CDP Runtime to CDP 7.1.9 CHF 2, and then you need to perform the following steps:

1. Open the Cloudera Manager Admin Console and go to the Knox service.
2. Click on the Knox Gateway Home URL.
3. Copy the YARN Resource Manager Web UI V2 URL from the Knox Gateway Home page.

For example, `https://knox-gateway.example.com:8443/gateway/cdp-proxy/yarnuiv2/`

4. Open the Cloudera Manager Admin Console and go to the YARN service.
5. Click the Configuration tab and search for resourcemanager\_config\_safety\_valve.
6. Add the Resource Manager Advanced Configuration Snippet (Safety Valve) for yarn-site.xml property, and specify its value by using the YARN Resource Manager Web UI V2 URL, copied earlier, as follows:

```
Name: yarn.web-proxy.gateway.url
Value: <YARN Resource Manager Web UI V2 URL>
```

7. Enter a Reason for Change and then click Save Changes.
8. Restart YARN.

### **Additional settings required for Spark after enabling HTTPS communication between the Resource Manager (RM) and the Application Master(AM)**

After enabling the HTTPS communication between the RM and AM, if yarn.resourcemanager.application-https.policy is set to STRICT, the following parameters must be specified for Spark:

```
spark.ssl.enabled=true
spark.ssl.protocol=TLSv1.2
```



```
spark.ssl.keystore=<pathtofile>
spark.ssl.keyStorePassword=<password>
--files /path/to/jks-file
```

When the `yarn.resourcemanager.application-https.policy` is set to `STRICT`, the tracking URL must be a HTTPS URL. If the tracking URL is not secure, use the below workaround to allow non-secure http connections between the RM and AM.

To allow non-https connections, set `yarn.resourcemanager.application-https.policy` to `LENIENT`. This allows the RM proxy to accept HTTP connections to AMs that opt not to use HTTPS.

## Unsupported Features

The following YARN features are currently not supported in Cloudera Data Platform:

- Application Timeline Server v2 (ATSv2)
- Auxiliary Services
- Container Resizing
- Distributed or Centralized Allocation of Opportunistic Containers
- Distributed Scheduling
- Docker on YARN (DockerContainerExecutor)
- 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 Zeppelin

Learn about the known issues in Zeppelin, the impact or changes to the functionality, and the workaround.

**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.

**BUG-125263: Zeppelin service move fails on clusters upgraded from HDP3.1.5**

Resolve the circular symlink issue on the Zeppelin node by linking the conf directory to a new directory under `/etc/zeppelin`:

- ```
# mkdir -p /etc/zeppelin/<version>/0
# rm /usr/hdp/<version>/zeppelin/conf
# ln -s /etc/zeppelin/<version>/0 /usr/hdp/<version>/zeppelin/conf
```

Where version is the HDP version. For example, `7.1.x-yyy`. Restart the Zeppelin server in Ambari.

CDPD-3090: Due to a configuration typo, functionality involving notebook repositories does not work

Due to a missing closing brace, access to the notebook repositories API is blocked by default.

From the CDP Management Console, go to Cloudera Manager for the cluster running Zeppelin. On the Zeppelin configuration page (Zeppelin serviceConfiguration), enter shiro urls in the Search field, and then add the missing closing brace to the notebook-repositories URL, as follows:

```
/api/notebook-repositories/** = authc, roles[{{zeppelin_admin_group}}]
```

Click Save Changes, and restart the Zeppelin service.

CDPD-2406: Logout button does not work

Clicking the Logout button in the Zeppelin UI logs you out, but then immediately logs you back in using SSO.

Close the browser.

OPSAPS-59802: Zeppelin and Livy roles should be co-located on the same host.

When installing or upgrading to CDP Private Cloud Base, you must co-locate all Zeppelin and Livy roles on the same cluster host due to an issue with certificate generation.

Known Issues in Apache ZooKeeper

Learn about the known issues in ZooKeeper, the impact or changes to the functionality, and the workaround.

OPSAPS-61188: Zookeeper start fails with custom user as contents inside /var/lib/zookeeper have "zookeeper" as owner instead of the custom user

In Cloudera Manager the Process Username for ZooKeeper can be changed from the default zookeeper value to any custom value. This configuration change in Cloudera Manager automatically changes the owner of the var/lib/zookeeper folder but keeps zookeeper as the owner of any folders or files inside var/lib/zookeeper, such as myid and version-2. As a result ZooKeeper fails to start because it needs to read the snapshots and txnlogs from the var/lib/zookeeper/version-2 folder when starting.

1. Ensure that you changed the Process Username to a username that exists on the OS.
2. Manually change the owner.
 - a. Log in to the node.
 - b. Recursively change the owner of var/lib/zookeeper using the chown -R command.

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.

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.zookeeper.ClientCnxnSocketNetty -Dzookeeper.ssl.keyStore.location=<PATH TO YOUR CONFIGURED KEYSTORE> -Dzookeeper.ssl.keyStore.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.secure=true" zookeeper-client -server <YOUR.ZOOKEEPER.SERVER-1>:2182,<YOUR.ZOOKEEPER.SERVER-2>:2182,<YOUR.ZOOKEEPER.SERVER-3>:2182
```

TLS v1.3 Support for CDP in Zookeeper

ZooKeeper does not support TLS v1.3.

To configure TLS v1.3, add the following properties in the Server Advanced Configuration Snippet (Safety Valve) for zoo.cfg file.

```
ssl.protocol=TLSv1.3
ssl.enabledProtocols=TLSv1.3,TLSv1.2
ssl.ciphersuites=TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256,TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256,TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384,TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384,TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256,TLS_ECDHE_RSA_WITH_CHACHA20_POLY1305_SHA256,TLS_DHE_RSA_WITH_AES_128_GCM_SHA256,TLS_DHE_RSA_WITH_AES_256_GCM_SHA384,TLS_AES_128_GCM_SHA256,TLS_AES_256_GCM_SHA384,TLS_CHACHA20_POLY1305_SHA256
```

Behavioral changes in Cloudera Runtime 7.1.9

You can review the changes in certain features or functionalities of components that have resulted in a change in behavior from the previously released version to this version of Cloudera Runtime 7.1.9.

Behavioral Changes in Cruise Control

Learn about the change in certain functionality of Cruise Control that has resulted in a change in behavior from the previously released version to this version of Cloudera Runtime.

Summary:

Instead of direct communication with ZooKeeper, Cruise Control uses the Kafka Reassignment API for partition management.

Previous behavior:

Cruise Control used Apache ZooKeeper for managing Kafka partitions.

New behavior:

Cruise Control uses Kafka Reassignment API for managing Kafka partitions.

Behavioral changes in Apache Hive

Learn about the change in certain functionality of Hive that has resulted in a change in behavior from the previously released version to this version of Cloudera Runtime.

Summary:

Change in the way dates are parsed from string by ignoring trailing invalid characters

Previous behavior:

[HIVE-20007](#) introduced changes in the way dates were parsed from strings. SQL functions or date operations involving invalid dates returned "null".

New behavior:

[HIVE-27586](#) extracts and returns a valid date from a string value if there is a valid date prefix in the string. This fix partially restores the behavior changes introduced as part of HIVE-20007 and also makes the current behavior of handling 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

Strong value	Behavior (before HIVE-20007)	Previous behavior (after HIVE-20007)	Current behavior (after HIVE-27586)
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.



Important: The behavior change introduced as part of this fix is only available from the CDP Private Cloud Base 7.1.9 version Cumulative hotfix 1.

Summary:

Handling invalid date formats in to_date function.

The behavior of the to_date function has changed between CDP Private Cloud Base versions 7.1.7 SP2 and 7.1.9 when handling invalid date formats.

Previous behavior:

Valid dates (e.g., YYYY-MM-DD): Returned correct results.

Invalid dates: Returned random, unexpected dates instead of NULL.

New behavior:

Following the changes introduced in [HIVE-28483](#):

Valid dates: Continue to return correct results.

Invalid dates (e.g., DD-MM-YYYY): Now return NULL.

Summary:

Change in the way date and timestamp values are parsed.

Previous behavior:

Some of the Hive date and timestamp functions, such as unix_timestamp(), from_unixtime(), date_format(), and cast() use the DateTimeFormatter class for printing and parsing date and timestamp objects. Prior to the CDP Private Cloud Base 7.1.7 SP2 version, these functions used the SimpleDateFormat class.



Note: If you are upgrading to CDP Private Cloud Base 7.1.9 from 7.1.7 SP1 or earlier releases, see the 7.1.7 SP2 or 7.1.8 release notes to understand the behavior changes related to date and timestamp functions that were introduced in these releases.

New behavior:

Starting from **CDP Private Cloud Base 7.1.9 version Cumulative hotfix 1**, a new configurable hive.datetime.formatter property is introduced through [HIVE-25576](#) that enables you to choose between SimpleDateFormat and DateTimeFormatter for the unix_timestamp and from_unixtime SQL functions.

Although the DateTimeFormatter class is an improvement over SimpleDateFormat, some users may want to retain the old behavior to ensure compatibility after migration, therefore, making it necessary for introducing this property.

The possible values for the `hive.datetime.formatter` property are 'DATETIME' and 'SIMPLE' representing `DateTimeFormatter` and `SimpleDateFormat` respectively. The default value is set to 'DATETIME'.

Summary:

Change in default value of the `hive.driver.parallel.compilation.global.limit` property

Previous behavior:

The default value for the `hive.driver.parallel.compilation.global.limit` property is set to "3".

New behavior:

The default value for the `hive.driver.parallel.compilation.global.limit` property is changed to "5", which helps in preventing queries from getting stuck because of a limit on the number of queries that can be compiled in parallel on a HiveServer (HS2) instance.

Summary:

Change in default value of the `hive.server2.tez.initialize.default.sessions` property

Previous behavior:

The default value for the `hive.server2.tez.initialize.default.sessions` property is set to "true"

New behavior:

The default value for the `hive.server2.tez.initialize.default.sessions` property is changed to "false" to prevent queries from waiting on the same Tez AM pool and thereby improving query performance.

If there are multiple queries running, you might notice that the queries are taking longer to complete because the default value for `hive.server2.tez.sessions.per.default.queue` is 1, which means only one query can run at a time. Therefore, depending on your resource availability and query concurrency/load on the server, you can set `hive.server2.tez.initialize.default.sessions` to "true" and increase the value of `hive.server2.tez.sessions.per.default.queue`.

Summary:

[HIVE-23100](#) introduces stricter type checking for comparison operations between columns and constants, affecting type combinations governed by the `hive.strict.checks.type.safety` property. This change ensures safer comparisons by enforcing stricter rules before any implicit type conversion occurs.

The table below outlines the behavior changes before and after HIVE-23100 for expressions of the form `E1 op E2` or `E2 op E1`, where `op` represents comparison operators such as `<`, `<=`, `=`, `>`, `>=`, `!=`, `E1` is a column reference, and `E2` is a constant.

Case	E1 Type	E2 Type	Before	After	Snippet
A	BIGINT	STRING	OK	ERROR/ WARN	<code>c_bigint = '9223372036854775807'</code>
B	DOUBLE	BIGINT	OK	WARN	<code>c_double = 9223372036854775807</code>

Previous behavior:

Hive implicitly converted the constant (E2) to the column's type (E1) before enforcing strict type checking, allowing queries to complete without errors or warnings. However, this approach carried risks of unintended results due to unsafe type conversions.

New behavior:

Strict type checking is performed before any implicit type conversion, and the outcome depends on the `hive.strict.checks.type.safety` property setting:

- If true, the query will throw a compilation error.

- If false, the query will execute but will generate a warning in the logs.

The error or warning will highlight that such comparisons are generally unsafe, with the message:

```
Unsafe compares between different types are disabled for safety reasons.
```



Note: HIVE-23100 does not focus on strict type conversions but has a much wider scope that aims to address various other shortcomings and problems in the compiler.

Behavioral Changes in Apache Kafka

Learn about the change in certain functionality of Kafka that has resulted in a change in behavior from the previously released version to this version of Cloudera Runtime.

Summary:

The Cluster Health Guarantee During Rolling Restart property is now set to healthy partitions stay healthy. This change is done so that a higher level of cluster health guarantees are provided by default.

Previous behavior:

The default value of the Cluster Health Guarantee During Rolling Restart property was set to none.

New behavior:

The default value of the Cluster Health Guarantee During Rolling Restart property is set to healthy partitions stay healthy.

Behavioral Changes in Ozone

Learn about the change in certain functionality of Ozone that has resulted in a change in behavior from the previously released version to this version of Cloudera Runtime.

Change in the behaviour of the Ozone admin pipeline list command.

In the CDP Private Cloud Base 7.1.9, running the Ozone admin pipeline list command displays the following output: Pipeline[Id: pipeline_id, Nodes: UUID1(hostname1/ip1)UUID2(hostname2/ip2)UUID3(hostname3/ip3), ReplicationConfig: RATIS/THREE, State:OPEN, leaderId:id, CreationTimestamp].

In the CDP Private Cloud Base 7.1.9 SP1, running the Ozone admin pipeline list command displays the following output: Pipeline[Id: pipeline_id, Nodes: UUID1(hostname1/ip1) ReplicationIndex: UUID2(hostname2/ip2) ReplicationIndex: UUID3(hostname3/ip3) ReplicationIndex: 0, ReplicationConfig: RATIS/THREE, State:OPEN, leaderId:id, CreationTimestamp].

Behavioral Changes in Apache Ranger

Learn about the change in certain functionality of Apache Ranger that has resulted in a change in behavior from the previously released version to this version of Cloudera Runtime.

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:

Added support for Ranger TagSync and UserSync HA

Previous behavior:

In versions < 7.1.9, port information for UserSync and Tagsync was not visible nor configurable.

New behavior:


In version 7.1.9 +, UserSync and TagSync port numbers are available at Cloudera Manager Ranger Configuration .

For example:

- 1. In Configuration Filters , highlight Ranger TagSync. Then, in Search, type http:


The following, configurable settings are now available:

Figure 1: Ranger TagSync port settings



2. In Configuration Filters , highlight Ranger UserSync. Then, in Search, type http:
- The following, configurable settings are now available:

Figure 2: Ranger UserSync port settings



For more information, see [Configuring Ranger UserSync and Tagsync High Availability](#).

Behavioral Changes in Schema Registry

Learn about the change in certain functionality of Schema Registry that has resulted in a change in behavior from the previously released version to this version of Cloudera Runtime.

Summary:

SchemaRegistryClient#runRetryableBlock is not thread safe

Previous behavior:

The default backoff policy for the Cloudera distributed Schema Registry Java client was `com.hortonworks.registries.schemaregistry.retry.policy.NOOPBackoffPolicy`. This backoff policy did not allow multiple retries even if more Schema Registry URLs were in use where multiple retries can occur.

New behavior:

The default backoff policy of the Schema Registry client has changed to `com.hortonworks.registries.schemaregistry.retry.policy.ExponentialBackoffPolicy`. This backoff policy allows retries of requests with the default configuration. If retries are not desirable in a client application, users need to explicitly set the `schema.registry.client.retry.policy.className` to `com.hortonworks.registries.schemaregistry.retry.policy.NOOPBackoffPolicy`.

Summary:

Increased SR server DB timeout for `JdbcStorageManager#nextId`

Previous behavior:

The Schema Registry database could easily timeout under heavy workload.

New behavior:

The database write timeout was increased from 3 seconds to 10 seconds. Additionally, the timeout related exception mapping is changed to be more informative and is retryable by the client.

Behavioral Changes in Apache Solr

Learn about the change in certain functionality of Apache Solr that has resulted in a change in behavior from the previously released version to this version of Cloudera Runtime.

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.

FIPS Compliant Changes in Apache Impala

As an administrator, you must understand the FIPS compliant changes in Impala before configuring Impala Web UI to diagnose issues with each daemon on a particular host, or perform other administrative actions such as cancelling a running query from the built-in web server's UI.

Cloudera Manager supports two methods of authentication for secure access to the Impala Catalog Server, Impala Daemon, and StateStore web servers: password-based authentication and SPNEGO authentication. From this release, Impala embedded Web Server will not support HTTP password-based authentication in FIPS approved mode since it's based on MD5 and does not comply with FIPS 140-2.

For details on FIPS encryption, see [Configure CDP with FIPS-compliant encryption](#).

Behavioral Changes in Apache Hadoop YARN

Learn about the change in certain functionality of YARN that has resulted in a change in behavior from the previously released version to this version of Cloudera Runtime.

Summary:

For YARN, mixed resource allocation mode allows you to specify the resources in mixed types.

Previous behavior:

Previously, mixed mode was not allowed and only a single type (Absolute value/Percentage value/Weight value) of allocation for all resources was allowed.

New behavior:

Now you can use a mix resources meaning different types for different resources.

CDP Private Cloud Base API Modifications and Removals

All the APIs that have either been modified or removed for this release of CDP Private Cloud Base.

CDP 7.1.7 SP2 and CDP 7.1.9 Components with API differences

The following is a list of all the APIs that have been modified or removed in this release of CDP Private Cloud Base. The APIs listed here are on a component by component basis and may include recommendations for how to manage the change.

API Compatibility changes in 7.1.9 for Hadoop

Removed or Modified APIs in CDP 7.1.9 for Hadoop and recommendations for how to handle them.

Apache Base Version of Hadoop in 7.1.7 SP2 was 3.1.1 and Apache Base Version of Hadoop in 7.1.9 is 3.1.1. The Cloudera version 7.1.9 has additional improvements over the Apache Base version.

Removed APIs in 7.1.9

The following APIs are no longer available for Hadoop in CDP 7.1.9

QueueInfo.newInstance

Method Removed

Package Name

org.apache.hadoop.yarn.api.records

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[YARN-10891](#)

Recommendation

Method signature changed, new parameter: maxParallelApps. If a client application calls this method, it should apply this new parameter as well.

Recompilation Required?

Yes

RemoteIterators.toArray

Method Removed

Package Name

org.apache.hadoop.util.functional

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[HADOOP-17511](#)

Recommendation

Method signature changed, new parameter: source. If a client application calls this method, it should apply this new parameter as well.

Recompilation Required?

Yes

Modified APIs in 7.1.9

The following APIs have been modified for Hadoop and include a description of the impact of the modification on their use.

QueueInfo

Abstract method void setMaxParallelApps (int) has been added to this class.

Package Name

org.apache.hadoop.yarn.api.records

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method setMaxParallelApps(int) in org.apache.hadoop.yarn.api.records.QueueInfo.

Reason for change

[YARN-10891](#)

Recommendation

If a class extends QueueInfo in a client applications, it must implement setMaxParallelApps abstract method.

Recompilation Required?

Yes

QueueInfo

Abstract method int getMaxParallelApps() has been added to this class.

Package Name

org.apache.hadoop.yarn.api.records

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method getMaxParallelApps() in org.apache.hadoop.yarn.api.records.QueueInfo.

Reason for change

[YARN-10891](#)

Recommendation

If a class extends QueueInfo in a client applications, it must implement getMaxParallelApps abstract method.

Recompilation Required?

Yes

QueueInfo

Abstract method int getMaxParallelApps() has been added to this class.

Package Name

org.apache.hadoop.yarn.api.records

Effect

A client program may be interrupted by AbstractMethodError exception.
Added abstract method is called in 2nd library version by the method org.apache.hadoop.yarn.client.cli.QueueCLI.printQueueInfo(PrintWriter; QueueInfo) and may not be implemented by old clients.

Reason for change

[YARN-10891](#)

Recommendation

If a class extends QueueInfo in a client applications, it must implement getMaxParallelApps abstract method.

Recompilation Required?

Yes

API Compatibility changes in 7.1.9 for Hive Warehouse Connector

Removed or Modified APIs in CDP 7.1.9 for Hive Warehouse Connector and recommendations for how to handle them.

Modified APIs in 7.1.9

The following APIs have been modified for Hive Warehouse Connector and include a description of the impact of the modification on their use.

CreateTableBuilder

Abstract method CreateTableBuilder database(String) has been added to this interface.

Package Name

com.hortonworks.hwc

Effect

Added method in the interface

Reason for change

Added support to specify customised database, fileformat and options for CreateTable API

Recommendation

In order to use customized database in createTable API, please recompile your application.

Recompilation Required?

Yes

CreateTableBuilder

Abstract method CreateTableBuilder fileFormat(String) has been added to this interface.

Package Name

com.hortonworks.hwc

Effect

Added method in the interface

Reason for change

Added support to specify customised database, fileformat and options for CreateTable API

Recommendation

In order to use customized file format in createTable API, please recompile your application.

Recompilation Required?

Yes

CreateTableBuilder

Abstract method CreateTableBuilder option(String; String) has been added to this interface.

Package Name

com.hortonworks.hwc

Effect

Added method in the interface

Reason for change

Added support to specify customised database, fileformat and options for CreateTable API

Recommendation

In order to use customized options in createTable API, please recompile your application.

Recompilation Required?

Yes

API Compatibility changes in 7.1.9 for Kafka

Removed or Modified APIs in CDP 7.1.9 for Kafka and recommendations for how to handle them.

Apache Base Version of Kafka in 7.1.7 SP2 was 2.5.0 and Apache Base Version of Kafka in 7.1.9 is 3.4.1

Removed APIs in 7.1.9

The following APIs are no longer available for Kafka in CDP 7.1.9

NoOffsetForPartitionException.partition

Method Removed

Package Name

org.apache.kafka.clients.consumer

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-12579: Remove various deprecated clients classes/methods for 3.0](#)

Recommendation

Deprecated. please use partitions

Recompilation Required?

Yes

Count.<init>

Method Removed

Package Name

org.apache.kafka.common.metrics.stats

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-8696: clean up Sum/Count/Total metrics](#)

Recommendation

Deprecated since 2.4 . Use WindowedCount instead

Recompilation Required?

Yes

Rate.SampledTotal.<init>

Method Removed

Package Name

org.apache.kafka.common.metrics.stats

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-8696: clean up Sum/Count/Total metrics](#)

Recommendation

Deprecated since 2.4 Use WindowedSum instead.

Recompilation Required?

Yes

Sum.<init>

Method Removed

Package Name

org.apache.kafka.common.metrics.stats

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-8696: clean up Sum/Count/Total metrics](#)

Recommendation

Deprecated since 2.4 Use WindowedSum instead.

Recompilation Required?

Yes

Total.<init>

Method Removed

Package Name

org.apache.kafka.common.metrics.stats

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-8696: clean up Sum/Count/Total metrics](#)

Recommendation

Deprecated since 2.4 Use CumulativeSum instead.

Recompilation Required?

Yes

Count

This class has been removed.

Package Name

org.apache.kafka.common.metrics.stats

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.metrics.stats.Count.

Reason for change

[KAFKA-8696: clean up Sum/Count/Total metrics](#)

Recommendation

Deprecated since 2.4 . Use WindowedCount instead

Recompilation Required?

Yes

Count

This class has been removed.

Package Name

org.apache.kafka.common.metrics.stats

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-8696: clean up Sum/Count/Total metrics](#)

Recommendation

Deprecated since 2.4 . Use WindowedCount instead

Recompilation Required?

Yes

Rate.SampledTotal

This class has been removed.

Package Name

org.apache.kafka.common.metrics.stats

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.metrics.stats.Rate.SampledTotal.

Reason for change

[KAFKA-8696: clean up Sum/Count/Total metrics](#)

Recommendation

Deprecated since 2.4 Use WindowedSum instead.

Recompilation Required?

Yes

Rate.SampledTotal

This class has been removed.

Package Name

org.apache.kafka.common.metrics.stats

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-8696: clean up Sum/Count/Total metrics](#)

Recommendation

Deprecated since 2.4 Use WindowedSum instead.

Recompilation Required?

Yes

Sum

This class has been removed.

Package Name

org.apache.kafka.common.metrics.stats

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.metrics.stats.Sum.

Reason for change

[KAFKA-8696: clean up Sum/Count/Total metrics](#)

Recommendation

Deprecated since 2.4 Use WindowedSum instead.

Recompilation Required?

Yes

Sum

This class has been removed.

Package Name

org.apache.kafka.common.metrics.stats

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-8696: clean up Sum/Count/Total metrics](#)

Recommendation

Deprecated since 2.4 Use WindowedSum instead.

Recompilation Required?

Yes

Total

This class has been removed.

Package Name

org.apache.kafka.common.metrics.stats

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.metrics.stats.Total.

Reason for change

[KAFKA-8696: clean up Sum/Count/Total metrics](#)

Recommendation

Deprecated since 2.4 Use CumulativeSum instead.

Recompilation Required?

Yes

Total

This class has been removed.

Package Name

org.apache.kafka.common.metrics.stats

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-8696: clean up Sum/Count/Total metrics](#)

Recommendation

Deprecated since 2.4 Use CumulativeSum instead.

Recompilation Required?

Yes

Modified APIs in 7.1.9

The following APIs have been modified for Kafka and include a description of the impact of the modification on their use.

KafkaStreams.close

Method Removed

Package Name

org.apache.kafka.streams

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-7477: Improve Streams close timeout semantics \(#5747\)](#)

Recommendation

Deprecated. Use close(Duration) instead; note, that close(Duration) has different semantics and does not block on zero, e.g., `Duration.ofMillis(0)`.

Recompilation Required?

Yes

KafkaStreams.metadataForKey

Method Removed

Package Name

org.apache.kafka.streams

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-6144: Add KeyQueryMetadata APIs to KafkaStreams](#)

Recommendation

Deprecated. Since 2.5. Use queryMetadataForKey(String, Object, Serializer) instead.

Recompilation Required?

Yes

KafkaStreams.store

Method Removed

Package Name

org.apache.kafka.streams

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-9487: Follow-up PR of Kafka-9445](#)

Recommendation

Deprecated. since 2.5 release; use store(StoreQueryParameters) instead

Recompilation Required?

Yes

StreamsBuilder.addGlobalStore

Method Removed

Package Name

org.apache.kafka.streams

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-10379: Implement the KIP-478 StreamBuilder#addGlobalStore\(\)](#)

Recommendation

Deprecated. use addGlobalStore(StoreBuilder, String, Consumed, ProcessorSupplier) instead

Recompilation Required?

Yes

StreamsConfig.getConsumerConfigs

Method Removed

Package Name

org.apache.kafka.streams

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-8284: enable static membership on KStream](#)

Recommendation

Deprecated. use getMainConsumerConfigs(String, String, int)

Recompilation Required?

Yes

StreamsMetrics.addLatencyAndThroughputSensor

Method Removed

Package Name

org.apache.kafka.streams

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-9230: Refactor user-customizable Streams metrics](#)**Recommendation**

Deprecated. since 2.5. Use addLatencyRateTotalSensor() instead.

Recompilation Required?

Yes

UnlimitedWindows.maintainMs

Method Removed

Package Name

org.apache.kafka.streams.kstream

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-7277: Migrate Streams API to Duration instead of longMs times](#)**Recommendation**

Deprecated since 2.1. Use Materialized.retention instead.

Recompilation Required?

Yes

UnlimitedWindows.startOn

Method Removed

Package Name

org.apache.kafka.streams.kstream

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-7277: Migrate Streams API to Duration instead of longMs times](#)**Recommendation**

Use UnlimitedWindows.startOn(Instant) instead

Recompilation Required?

Yes

UnlimitedWindows.until

Method Removed

Package Name

org.apache.kafka.streams.kstream

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-7277: Migrate Streams API to Duration instead of longMs times](#)

Recommendation

Deprecated since 2.1.

Recompilation Required?

Yes

Windows.maintainMs

Method Removed

Package Name

org.apache.kafka.streams.kstream

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-7277: Migrate Streams API to Duration instead of longMs times](#)

Recommendation

Deprecated since 2.1. This function should not be used anymore, since JoinWindows.until(long) is deprecated in favor of JoinWindows.grace(Duration).

Recompilation Required?

Yes

Windows.segments

Method Removed

Package Name

org.apache.kafka.streams.kstream

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-7277: Migrate Streams API to Duration instead of longMs times](#)

Recommendation

Deprecated since 2.1 Override segmentInterval() instead.

Recompilation Required?

Yes

Windows.until

Method Removed

Package Name

org.apache.kafka.streams.kstream

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-7277: Migrate Streams API to Duration instead of longMs times](#)

Recommendation

Deprecated since 2.1. Use Materialized.withRetention(Duration) or directly configure the retention in a store supplier and use Materialized.as(WindowBytesStoreSupplier).

Recompilation Required?

Yes

Serialized<K;V>

This class has been removed.

Package Name

org.apache.kafka.streams.kstream

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.streams.kstream.SerializedK;V.

Reason for change

[KAFKA-7406: Name join group repartition topics](#)

Recommendation

Deprecated since 2.1. Use Grouped instead

Recompilation Required?

Yes

Serialized<K;V>

This class has been removed.

Package Name

org.apache.kafka.streams.kstream

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-7406: Name join group repartition topics](#)

Recommendation

Deprecated since 2.1. Use Grouped instead

Recompilation Required?

Yes

Windows<W>

Field segments of type int has been removed from this class.

Package Name

org.apache.kafka.streams.kstream

Effect

Recompilation of a client program may be terminated with the message: cannot find variable segments in org.apache.kafka.streams.kstream.WindowsW.

Reason for change

[KAFKA-7080: replace numSegments with segmentInterval](#)

Recommendation

Deprecated since 2.1 Override segmentInterval() instead.

Recompilation Required?

Yes

Windows<W>

Field segments of type int has been removed from this class.

Package Name

org.apache.kafka.streams.kstream

Effect

A client program may be interrupted by NoSuchFieldError exception.

Reason for change

[KAFKA-7080: replace numSegments with segmentInterval](#)

Recommendation

Deprecated since 2.1 Override segmentInterval() instead.

Recompilation Required?

Yes

DefaultPartitionGrouper.<init>

Method Removed

Package Name

org.apache.kafka.streams.processor

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-8927: Deprecate PartitionGrouper interface](#)

Recommendation

Deprecated since 2.4 release; will be removed in 3.0.0 via KAFKA-7785

Recompilation Required?

Yes

DefaultPartitionGrouper.maxNumPartitions

Method Removed

Package Name

org.apache.kafka.streams.processor

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-8927: Deprecate PartitionGrouper interface](#)

Recommendation

Deprecated since 2.4 release; will be removed in 3.0.0 via KAFKA-7785

Recompilation Required?

Yes

DefaultPartitionGrouper.partitionGroups

Method Removed

Package Name

org.apache.kafka.streams.processor

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-8927: Deprecate PartitionGrouper interface](#)

Recommendation

Deprecated since 2.4 release; will be removed in 3.0.0 via KAFKA-7785

Recompilation Required?

Yes

MockProcessorContext.forward

Method Removed

Package Name

org.apache.kafka.streams.processor

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

MINOR: cleanup deprecation annotations

Recommendation

please use ProcessorContext.forward(Object, Object, To) instead

Recompilation Required?

Yes

MockProcessorContext.schedule

Method Removed

Package Name

org.apache.kafka.streams.processor

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-7277: Migrate Streams API to Duration instead of longMs times](#)

Recommendation

Deprecated since 2.1. Use WindowBytesStoreSupplier.segmentIntervalMs() instead.

Recompilation Required?

Yes

PartitionGrouper.partitionGroups

Method Removed

Package Name

org.apache.kafka.streams.processor

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-8927: Deprecate PartitionGrouper interface](#)

Recommendation

Deprecated since 2.4 release; will be removed in 3.0.0 via KAFKA-7785

Recompilation Required?

Yes

ProcessorContext.forward

Method Removed

Package Name

org.apache.kafka.streams.processor

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-6454: Allow timestamp manipulation in Processor API](#)

Recommendation

please use ProcessorContext.forward(Object, Object, To) instead

Recompilation Required?

Yes

ProcessorContext.schedule

Method Removed

Package Name

org.apache.kafka.streams.processor

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-6454: Allow timestamp manipulation in Processor API](#)**Recommendation**

Use ProcessorContext.schedule(Duration, PunctuationType, Punctuator) instead

Recompilation Required?

Yes

UsePreviousTimeOnInvalidTimestamp.<init>

Method Removed

Package Name

org.apache.kafka.streams.processor

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-8953: Rename UsePreviousTimeOnInvalidTimestamp to UsePartitionTimeOnInvalidTimestamp](#)**Recommendation**

Deprecated since 2.5. Use UsePartitionTimeOnInvalidTimestamp instead

Recompilation Required?

Yes

UsePreviousTimeOnInvalidTimestamp.extract

Method Removed

Package Name

org.apache.kafka.streams.processor

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-8953: Rename UsePreviousTimeOnInvalidTimestamp to UsePartitionTimeOnInvalidTimestamp](#)**Recommendation**

Deprecated since 2.5. Use UsePartitionTimeOnInvalidTimestamp instead

Recompilation Required?

Yes

UsePreviousTimeOnInvalidTimestamp.onInvalidTimestamp

Method Removed

Package Name

org.apache.kafka.streams.processor

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-8953: Rename UsePreviousTimeOnInvalidTimestamp to UsePartitionTimeOnInvalidTimestamp](#)

Recommendation

Deprecated since 2.5. Use UsePartitionTimeOnInvalidTimestamp instead

Recompilation Required?

Yes

DefaultPartitionGrouper

This class has been removed.

Package Name

org.apache.kafka.streams.processor

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.streams.processor.DefaultPartitionGrouper.

Reason for change

[KAFKA-8927: Deprecate PartitionGrouper interface](#)

Recommendation

Deprecated since 2.4 release; will be removed in 3.0.0 via KAFKA-7785

Recompilation Required?

Yes

DefaultPartitionGrouper

This class has been removed.

Package Name

org.apache.kafka.streams.processor

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-8927: Deprecate PartitionGrouper interface](#)

Recommendation

Deprecated since 2.4 release; will be removed in 3.0.0 via KAFKA-7785

Recompilation Required?

Yes

PartitionGrouper

This interface has been removed.

Package Name

org.apache.kafka.streams.processor

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.streams.processor.PartitionGrouper.

Reason for change

[KAFKA-8927: Deprecate PartitionGrouper interface](#)

Recommendation

Deprecated since 2.4 release; will be removed in 3.0.0 via KAFKA-7785

Recompilation Required?

Yes

PartitionGrouper

This interface has been removed.

Package Name

org.apache.kafka.streams.processor

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-8927: Deprecate PartitionGrouper interface](#)

Recommendation

Deprecated since 2.4 release; will be removed in 3.0.0 via KAFKA-7785

Recompilation Required?

Yes

UsePreviousTimeOnInvalidTimestamp

This class has been removed.

Package Name

org.apache.kafka.streams.processor

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.streams.processor.UsePreviousTimeOnInvalidTimestamp.

Reason for change

[KAFKA-8953: Rename UsePreviousTimeOnInvalidTimestamp to UsePartitionTimeOnInvalidTimestamp](#)

Recommendation

Deprecated since 2.5. Use UsePartitionTimeOnInvalidTimestamp instead

Recompilation Required?

Yes

UsePreviousTimeOnInvalidTimestamp

This class has been removed.

Package Name

org.apache.kafka.streams.processor

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-8953: Rename UsePreviousTimeOnInvalidTimestamp to UsePartitionTimeOnInvalidTimestamp](#)

Recommendation

Deprecated since 2.5. Use UsePartitionTimeOnInvalidTimestamp instead

Recompilation Required?

Yes

ReadOnlyWindowStore.fetch

Method Removed

Package Name

org.apache.kafka.streams.state

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-7277: Migrate Streams API to Duration instead of longMs times](#)

Recommendation

Use `ReadOnlyWindowStore.fetch(Object, Object, Instant, Instant)` instead

Recompilation Required?

Yes

ReadOnlyWindowStore.fetchAll

Method Removed

Package Name

`org.apache.kafka.streams.state`

Effect

A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change

[KAFKA-7277: Migrate Streams API to Duration instead of longMs times](#)

Recommendation

Use `ReadOnlyWindowStore.fetch(Object, Object, Instant, Instant)` instead

Recompilation Required?

Yes

Stores.persistentSessionStore

Method Removed

Package Name

`org.apache.kafka.streams.state`

Effect

A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change

[KAFKA-7277: Migrate Streams API to Duration instead of longMs times](#)

Recommendation

Deprecated since 2.1 Use `Stores.persistentSessionStore(String, Duration)` instead

Recompilation Required?

Yes

Stores.persistentWindowStore

Method Removed

Package Name

`org.apache.kafka.streams.state`

Effect

A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change

[KAFKA-7277: Migrate Streams API to Duration instead of longMs times](#)

Recommendation

Deprecated since 2.1 Use `Stores.persistentWindowStore(String, Duration, Duration, boolean)` instead

Recompilation Required?

Yes

WindowBytesStoreSupplier.segments

Method Removed

Package Name

org.apache.kafka.streams.state

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-7080: replace numSegments with segmentInterval](#)**Recommendation**

Deprecated since 2.1. Use WindowBytesStoreSupplier.segmentIntervalMs() instead.

Recompilation Required?

Yes

WindowStore.put

Method Removed

Package Name

org.apache.kafka.streams.state

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-7245: Deprecate WindowStore#put\(key, value\)](#)**Recommendation**

Deprecated as timestamp is not provided for the key-value pair, this causes inconsistency to identify the window frame to which the key belongs. Use WindowStore.put(Object, Object, long) instead.

Recompilation Required?

Yes

API Compatibility changes in 7.1.9 for Oozie

Removed or Modified APIs in CDP 7.1.9 for Oozie and recommendations for how to handle them.

Apache Base Version of Oozie in 7.1.7 SP2 was 5.1.0 and Apache Base Version of Oozie in 7.1.9 is 5.1.0. The Cloudera version 7.1.9 has additional improvements over the Apache Base version.

Modified APIs in 7.1.9

The following APIs have been modified for Oozie and include a description of the impact of the modification on their use.

ActionExecutor.Context

Abstract method String[] getCallbackUrls(String) has been added to this interface.

Package Name

org.apache.oozie.action

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method getCallbackUrls(String) in org.apache.oozie.action.ActionExecutor.Context.

Reason for change

Ability for Oozie to not rely on its LoadBalancer internally

Recommendation

Use the same implementation as ActionExecutorContext or return a single url wrapped into an array.

Recompilation Required?

Yes

ActionExecutor.Context

Abstract method `FileSystem getFileSystemForDir(Path)` has been added to this interface.

Package Name

`org.apache.oozie.action`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getFileSystemForDir(Path)` in `org.apache.oozie.action.ActionExecutor.Context`.

Reason for change

In order to make action dir base path configurable and Cloud FS ready

Recommendation

Use a similar implementation as `ActionExecutorContext`.

Recompilation Required?

Yes

JavaActionExecutor.setCredentialTokens

Access level has been changed from protected to package-private.

Package Name

`org.apache.oozie.action.hadoop`

Effect

Recompilation of a client program may be terminated with the message: `setCredentialTokens(Credentials; Configuration; ActionExecutor.Context; WorkflowAction; MapString; CredentialsProperties)` has package-private access in `org.apache.oozie.action.hadoop.JavaActionExecutor`.

Reason for change

Oozie to handle file-system credentials coming from various places

Recommendation

Users should not invoke this method in their own `ActionExecutor` implementation. If they absolutely need to invoke this, then their own `ActionExecutor` implementation should be in the `org.apache.oozie.action.hadoop` package.

Recompilation Required?

Yes

API Compatibility changes in 7.1.9 for ORC

Removed or Modified APIs in CDP 7.1.9 for ORC and recommendations for how to handle them.

Apache Base Version of ORC in 7.1.7 SP2 was 1.5.1 and Apache Base Version of ORC in 7.1.9 is 1.5.1. The Cloudera version 7.1.9 has additional improvements over the Apache Base version.

Removed APIs in 7.1.9

The following APIs are no longer available for ORC in CDP 7.1.9

OrcUtils.appendOrcTypesRebuildSubtypes

Method Removed

Package Name

`org.apache.orc`

Effect

A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change

[ORC-522: Add user type annotations. Was removed as part of column level encryption changes to ORC](#)

Recommendation

use the method `appendOrcTypes(List<OrcProto.Type> result, TypeDescription typeDescr)` instead

Recompilation Required?

Yes

API Compatibility changes in 7.1.9 for Ozone

Removed or Modified APIs in CDP 7.1.9 for Ozone and recommendations for how to handle them.

Apache Base Version of Ozone in 7.1.7 SP2 was 1.1.0 and Apache Base Version of Ozone in 7.1.9 is 1.2.0

Removed APIs in 7.1.9

The following APIs are no longer available for Ozone in CDP 7.1.9

[OzoneBucket.addAcls](#)

Method Removed

Package Name

`org.apache.hadoop.ozone.client`

Effect

A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change

[HDDS-4585. Support bucket acl operation in S3g](#)

Recommendation

Method ``addAcls`` is renamed to ``addAcl``.

Recompilation Required?

Yes

[OzoneBucket.removeAcls](#)

Method Removed

Package Name

`org.apache.hadoop.ozone.client`

Effect

A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change

[HDDS-4585. Support bucket acl operation in S3g](#)

Recommendation

Method ``removeAcls`` is renamed to ``removeAcl``.

Recompilation Required?

Yes

API Compatibility changes in 7.1.9 for Spark

Removed or Modified APIs in CDP 7.1.9 for Spark and recommendations for how to handle them.

Apache Base Version of Spark in 7.1.7 SP2 was 2.4.7 and Apache Base Version of Spark in 7.1.9 is 2.4.8

Removed APIs in 7.1.9

The following APIs are no longer available for Spark in CDP 7.1.9

[UnsafeInMemorySorter.free](#)

Method Removed

Package Name

`org.apache.spark.util.collection.unsafe.sort`

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[SPARK-32901](#)

Recommendation

Use UnsafeInMemorySorter.freeMemory method instead of UnsafeInMemorySorter.free.

Recompilation Required?

Yes

UnsafeInMemorySorter.reset

Method Removed

Package Name

org.apache.spark.util.collection.unsafe.sort

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[SPARK-32901](#)

Recommendation

Method UnsafeInMemorySorter.reset() removed. Please see its code for alternatives.

Recompilation Required?

Yes

ShuffleIndexInformation.getSize

Method Removed

Package Name

org.apache.spark.network.shuffle

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[SPARK-33206](#)

Recommendation

Method getSize removed. Use getRetainedMemorySize() instead.

Recompilation Required?

Yes

CompactibleFileStreamLog.compactLogs

Method Removed

Package Name

org.apache.spark.sql.execution.streaming

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[SPARK-30462](#)

Recommendation

Method compactLogs has been removed. Use shouldRetain(log: T) instead

Recompilation Required?

Yes

FileStreamSinkLog.compactLogs

Method Removed

Package Name

org.apache.spark.sql.execution.streaming

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[SPARK-30462](#)**Recommendation**

Method compactLogs has been removed. Use shouldRetain(log: T) instead

Recompilation Required?

Yes

FileStreamSourceLog.compactLogs

Method Removed

Package Name

org.apache.spark.sql.execution.streaming

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[SPARK-30462](#)**Recommendation**

Method compactLogs has been removed. Use shouldRetain(log: T) instead

Recompilation Required?

Yes

Modified APIs in 7.1.9

The following APIs have been modified for Spark and include a description of the impact of the modification on their use.

UnsafeSorterIterator

Abstract method long getCurrentPageNumber() has been added to this class.

Package Name

org.apache.spark.util.collection.unsafe.sort

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method getCurrentPageNumber() in org.apache.spark.util.collection.unsafe.sort.UnsafeSorterIterator.

Reason for change[SPARK-32900](#)**Recommendation**

New abstract method getCurrentPageNumber added to UnsafeSorterIterator. Please recompile your application if you want to use this.

Recompilation Required?

Yes

CompactibleFileStreamLog<T>

Abstract method SeqT compactLogs(SeqT) has been removed from this class.

Package Name

org.apache.spark.sql.execution.streaming

Effect

Recompilation of a client program may be terminated with the message: cannot find method compactLogs(SeqT) in class org.apache.spark.sql.execution.streaming.CompactibleFileStreamLogT.

Reason for change

[SPARK-30462](#)

Recommendation

Method compactLogs has been removed. Use shouldRetain(log: T) instead

Recompilation Required?

Yes

CompactibleFileStreamLog<T>

Abstract method SeqT compactLogs(SeqT) has been removed from this class.

Package Name

org.apache.spark.sql.execution.streaming

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[SPARK-30462](#)

Recommendation

Method compactLogs has been removed. Use shouldRetain(log: T) instead

Recompilation Required?

Yes

CDP 7.1.7 SP2 and CDP 7.1.9 Compatible Components

The Java APIs for the components mentioned are 100% backward compatible between CDP Private Cloud Base 7.1.7 SP2 and CDP Private Cloud Base 7.1.9.

- Arrow
- Atlas
- Avro
- HBase
- Hive
- Knox
- Kudu
- Parquet
- Search
- Solr
- Spark_Schema_Registry
- Spark-Solr
- Zeppelin
- Zookeeper

CDP 7.1.8 CHF 12 and CDP 7.1.9 Components with API differences

The following is a list of all the APIs that have been modified or removed in this release of CDP Private Cloud Base. The APIs listed here are on a component by component basis and may include recommendations for how to manage the change.

API Compatibility changes in 7.1.9 for Oozie

Removed or Modified APIs in CDP 7.1.9 for Oozie and recommendations for how to handle them.

Apache Base Version of Oozie in 7.1.8 was 5.1.0 and Apache Base Version of Oozie in 7.1.9 is 5.1.0. The Cloudera version 7.1.9 has additional improvements over the Apache Base version.

Modified APIs in 7.1.9

The following APIs have been modified for Oozie and include a description of the impact of the modification on their use.

ActionExecutor.Context

Abstract method `String[] getCallbackUrls(String)` has been added to this interface.

Package Name

`org.apache.oozie.action`

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getCallbackUrls(String)` in `org.apache.oozie.action.ActionExecutor.Context`.

Reason for change

Ability for Oozie to not rely on its LoadBalancer internally

Recommendation

Use the same implementation as `ActionExecutorContext` or return a single url wrapped into an array.

Recompilation Required?

Yes

API Compatibility changes in 7.1.9 for Kafka

Removed or Modified APIs in CDP 7.1.9 for Kafka and recommendations for how to handle them.

Apache Base Version of Kafka in 7.1.8 was 3.1.1 and Apache Base Version of Kafka in 7.1.9 is 3.4.1

Removed APIs in 7.1.9

The following APIs are no longer available for Kafka in CDP 7.1.9

AccessTokenRetriever.close

Method Removed

Package Name

`org.apache.kafka.common.security.oauthbearer.secured`

Effect

A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the `org.apache.kafka.common.security.oauthbearer.internals.secured` package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

AccessTokenRetriever.retrieve

Method Removed

Package Name

`org.apache.kafka.common.security.oauthbearer.secured`

Effect

A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

AccessTokenRetrieverFactory.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

AccessTokenRetrieverFactory.create

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

AccessTokenValidator.validate

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

AccessTokenValidatorFactory.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

AccessTokenValidatorFactory.create

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

BasicOAuthBearerToken.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

BasicOAuthBearerToken.lifetimeMs

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

BasicOAuthBearerToken.principalName

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

BasicOAuthBearerToken.scope

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

BasicOAuthBearerToken.startTimeMs

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

BasicOAuthBearerToken.value

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ValidatorAccessTokenValidator.ClaimSupplier.get

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ClaimValidationUtils.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ClaimValidationUtils.validateClaimNameOverride

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ClaimValidationUtils.validateExpiration

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ClaimValidationUtils.validateIssuedAt

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ClaimValidationUtils.validateScopes

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ClaimValidationUtils.validateSubject

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

CloseableVerificationKeyResolver.close

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ConfigurationUtils.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ConfigurationUtils.get

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ConfigurationUtils.validateFile

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ConfigurationUtils.validateInteger

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ConfigurationUtils.validateLong

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ConfigurationUtils.validateString

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ConfigurationUtils.validateUrl

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

FileTokenRetriever.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

FileTokenRetriever.init

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

FileTokenRetriever.retrieve

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

HttpAccessTokenRetriever.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

HttpAccessTokenRetriever.post

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

[HttpAccessTokenRetriever.retrieve](#)

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

[Initable.init](#)

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

[JaasOptionsUtils.createSSLSocketFactory](#)

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

JaasOptionsUtils.getOptions

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

JaasOptionsUtils.getSslClientConfig

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

JaasOptionsUtils.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

JaasOptionsUtils.shouldCreateSSLSocketFactory

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

JaasOptionsUtils.validateString

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

JwksFileVerificationKeyResolver.init

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

JwksFileVerificationKeyResolver.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

JwksFileVerificationKeyResolver.resolveKey

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

LoginAccessTokenValidator.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

LoginAccessTokenValidator.validate

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

RefreshingHttpsJwks.close

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

RefreshingHttpsJwks.getJsonWebKeys

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

RefreshingHttpsJwks.getLocation

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

RefreshingHttpsJwks.init

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

RefreshingHttpsJwks.maybeExpediteRefresh

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

RefreshingHttpsJwks.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

RefreshingHttpsJwksVerificationKeyResolver.close

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

RefreshingHttpsJwksVerificationKeyResolver.init

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

RefreshingHttpsJwksVerificationKeyResolver.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

RefreshingHttpsJwksVerificationKeyResolver.resolveKey

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

Retry.execute

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

Retry.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

Retryable.call

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

SerializedJwt.getHeader

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

SerializedJwt.getPayload

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

SerializedJwt.getSignature

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

SerializedJwt.getToken

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

SerializedJwt.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

UnretryableException.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ValidateException.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ValidatorAccessTokenValidator.validate

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ValidatorAccessTokenValidator.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

VerificationKeyResolverFactory.create

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

VerificationKeyResolverFactory.<init>

Method Removed

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

KStream.process

Method Removed

Package Name

org.apache.kafka.streams.kstream

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

[KAFKA-13654: Extend KStream process with new Processor API](#)

Recommendation

Deprecated Since 3.0. Use
`KStream.process(org.apache.kafka.streams.processor.api.ProcessorSupplier,
 org.apache.kafka.streams.kstream.Named, java.lang.String...)` instead.

Recompilation Required?

Yes

KStream.process

Method Removed

Package Name

`org.apache.kafka.streams.scala.kstream`

Effect

A client program may be interrupted by `NoSuchMethodError` exception.

Reason for change

[KAFKA-13654: Extend KStream process with new Processor API](#)

Recommendation

Deprecated Since 3.0. Use
`KStream.process(org.apache.kafka.streams.processor.api.ProcessorSupplier,
 org.apache.kafka.streams.kstream.Named, java.lang.String...)` instead.

Recompilation Required?

Yes

Modified APIs in 7.1.9

The following APIs have been modified for Kafka and include a description of the impact of the modification on their use.

AccessTokenRetrieverFactory

This class has been removed.

Package Name

`org.apache.kafka.common.security.oauthbearer.secured`

Effect

Recompilation of a client program may be terminated with the message: cannot find class
`org.apache.kafka.common.security.oauthbearer.secured.AccessTokenRetrieverFactory`.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the `org.apache.kafka.common.security.oauthbearer.internals.secured` package, use
 this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

AccessTokenValidatorFactory

This class has been removed.

Package Name

`org.apache.kafka.common.security.oauthbearer.secured`

Effect

Recompilation of a client program may be terminated with the message: cannot find class
`org.apache.kafka.common.security.oauthbearer.secured.AccessTokenValidatorFactory`.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

BasicOAuthBearerToken

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.BasicOAuthBearerToken.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ClaimValidationUtils

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.ClaimValidationUtils.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ConfigurationUtils

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.ConfigurationUtils.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

FileTokenRetriever

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.FileTokenRetriever.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

HttpAccessTokenRetriever

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.HttpAccessTokenRetriever.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

JaasOptionsUtils

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.JaasOptionsUtils.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

JwksFileVerificationKeyResolver

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.JwksFileVerificationKeyResolver.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

LoginAccessTokenValidator

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.LoginAccessTokenValidator.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

OAuthBearerLoginCallbackHandler

Removed super-interface org.apache.kafka.common.security.auth.AuthenticateCallbackHandler.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find method in class org.apache.kafka.common.security.oauthbearer.secured.OAuthBearerLoginCallbackHandler.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

OAuthBearerLoginCallbackHandler

Field CLIENT_ID_CONFIG of type java.lang.String with the compile-time constant value "clientId" has been removed from this class.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find variable CLIENT_ID_CONFIG in org.apache.kafka.common.security.oauthbearer.secured.OAuthBearerLoginCallbackHandler.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

OAuthBearerLoginCallbackHandler

Field CLIENT_ID_DOC of type java.lang.String with the compile-time constant value "The OAuth/OIDC identity provider-issued client ID to uniquely identify the service account to use for authentication for this client. The value must be paired with a corresponding clientSecret value and is provided to the OAuth provider using the OAuth clientcredentials grant type." has been removed from this class.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find variable CLIENT_ID_DOC in org.apache.kafka.common.security.oauthbearer.secured.OAuthBearerLoginCallbackHandler.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

OAuthBearerLoginCallbackHandler

Field CLIENT_SECRET_CONFIG of type java.lang.String with the compile-time constant value "clientSecret" has been removed from this class.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find variable CLIENT_SECRET_CONFIG in org.apache.kafka.common.security.oauthbearer.secured.OAuthBearerLoginCallbackHandler.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

OAuthBearerLoginCallbackHandler

Field CLIENT_SECRET_DOC of type java.lang.String with the compile-time constant value "The OAuth/OIDC identity provider-issued client secret serves as a similar function as a password to the clientId account and identifies the service account to use for authentication for this client. The value must be paired with a corresponding clientId value and is provided to the OAuth provider using the OAuth clientcredentials grant type." has been removed from this class.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find variable CLIENT_SECRET_DOC in org.apache.kafka.common.security.oauthbearer.secured.OAuthBearerLoginCallbackHandler.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

OAuthBearerLoginCallbackHandler

Field SCOPE_CONFIG of type java.lang.String with the compile-time constant value "scope" has been removed from this class.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find variable SCOPE_CONFIG in org.apache.kafka.common.security.oauthbearer.secured.OAuthBearerLoginCallbackHandler.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

OAuthBearerLoginCallbackHandler

Field SCOPE_DOC of type java.lang.String with the compile-time constant value "The (optional) HTTP/HTTPS login request to the token endpoint (sas.oauthbearer.token.endpoint.url) may need to specify an OAuth \"scope\". If so; the scope is used to provide the value to include with the login request." has been removed from this class.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find variable SCOPE_DOC in org.apache.kafka.common.security.oauthbearer.secured.OAuthBearerLoginCallbackHandler.

Reason for change

[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

OAuthBearerValidatorCallbackHandler

Removed super-interface org.apache.kafka.common.security.auth.AuthenticateCallbackHandler.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find method in class org.apache.kafka.common.security.oauthbearer.secured.OAuthBearerValidatorCallbackHandler.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

RefreshingHttpsJwks

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.RefreshingHttpsJwks.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

RefreshingHttpsJwksVerificationKeyResolver

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.RefreshingHttpsJwksVerificationKeyResolver.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

Retry<R>

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.RetryR.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

SerializedJwt

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.SerializedJwt.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

UnretryableException

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.UnretryableException.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ValidateException

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.ValidateException.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ValidatorAccessTokenValidator

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.ValidatorAccessTokenValidator.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

VerificationKeyResolverFactory

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.VerificationKeyResolverFactory.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

AccessTokenRetriever

This interface has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.AccessTokenRetriever.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

AccessTokenValidator

This interface has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.AccessTokenValidator.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

CloseableVerificationKeyResolver

This interface has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.CloseableVerificationKeyResolver.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

Initable

This interface has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.Initable.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

Retryable<R>

This interface has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.RetryableR.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ValidatorAccessTokenValidator.ClaimSupplier<T>

This interface has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

Recompilation of a client program may be terminated with the message: cannot find class org.apache.kafka.common.security.oauthbearer.secured.ValidatorAccessTokenValidator.ClaimSupplierT.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

AccessTokenRetrieverFactory

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

AccessTokenValidatorFactory

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

BasicOAuthBearerToken

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ClaimValidationUtils

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ConfigurationUtils

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

FileTokenRetriever

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

HttpAccessTokenRetriever

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

JaasOptionsUtils

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

JwksFileVerificationKeyResolver

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

LoginAccessTokenValidator

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

OAuthBearerLoginCallbackHandler

Removed super-interface org.apache.kafka.common.security.auth.AuthenticateCallbackHandler.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

OAuthBearerValidatorCallbackHandler

Removed super-interface org.apache.kafka.common.security.auth.AuthenticateCallbackHandler.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

RefreshingHttpsJwks

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

RefreshingHttpsJwksVerificationKeyResolver

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

Retry<R>

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by `NoClassDefFoundError` exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the `org.apache.kafka.common.security.oauthbearer.internals.secured` package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

SerializedJwt

This class has been removed.

Package Name

`org.apache.kafka.common.security.oauthbearer.secured`

Effect

A client program may be interrupted by `NoClassDefFoundError` exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the `org.apache.kafka.common.security.oauthbearer.internals.secured` package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

UnretryableException

This class has been removed.

Package Name

`org.apache.kafka.common.security.oauthbearer.secured`

Effect

A client program may be interrupted by `NoClassDefFoundError` exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the `org.apache.kafka.common.security.oauthbearer.internals.secured` package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ValidateException

This class has been removed.

Package Name

`org.apache.kafka.common.security.oauthbearer.secured`

Effect

A client program may be interrupted by `NoClassDefFoundError` exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the `org.apache.kafka.common.security.oauthbearer.internals.secured` package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ValidatorAccessTokenValidator

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

VerificationKeyResolverFactory

This class has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

AccessTokenRetriever

This interface has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by NoClassDefFoundError exception.

Reason for change[KAFKA-13725](#)**Recommendation**

Class moved to the org.apache.kafka.common.security.oauthbearer.internals.secured package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

AccessTokenValidator

This interface has been removed.

Package Name

org.apache.kafka.common.security.oauthbearer.secured

Effect

A client program may be interrupted by `NoClassDefFoundError` exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the `org.apache.kafka.common.security.oauthbearer.internals.secured` package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

CloseableVerificationKeyResolver

This interface has been removed.

Package Name

`org.apache.kafka.common.security.oauthbearer.secured`

Effect

A client program may be interrupted by `NoClassDefFoundError` exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the `org.apache.kafka.common.security.oauthbearer.internals.secured` package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

Initable

This interface has been removed.

Package Name

`org.apache.kafka.common.security.oauthbearer.secured`

Effect

A client program may be interrupted by `NoClassDefFoundError` exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the `org.apache.kafka.common.security.oauthbearer.internals.secured` package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

Retryable<R>

This interface has been removed.

Package Name

`org.apache.kafka.common.security.oauthbearer.secured`

Effect

A client program may be interrupted by `NoClassDefFoundError` exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the `org.apache.kafka.common.security.oauthbearer.internals.secured` package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

ValidatorAccessTokenValidator.ClaimSupplier<T>

This interface has been removed.

Package Name

`org.apache.kafka.common.security.oauthbearer.secured`

Effect

A client program may be interrupted by `NoClassDefFoundError` exception.

Reason for change

[KAFKA-13725](#)

Recommendation

Class moved to the `org.apache.kafka.common.security.oauthbearer.internals.secured` package, use this class instead (KAFKA-13725: KIP-768)

Recompilation Required?

Yes

CDP 7.1.8 CHF 12 and CDP 7.1.9 Compatible Components

The Java APIs for the components mentioned are 100% backward compatible between CDP Private Cloud Base 7.1.8 and CDP Private Cloud Base 7.1.9.

- Arrow
- Atlas
- Avro
- Hadoop
- HBase
- HBase_MCC
- Hive
- Hive Warehouse Connector
- Knox
- Kudu
- ORC
- Ozone
- Parquet
- Search
- Solr
- Spark
- Spark_Schema_Registry
- Spark-Solr
- Zeppelin
- Zookeeper

CDP Private Cloud Base REST API Modifications and Removals

All the REST APIs that have either been modified or removed for this release of CDP Private Cloud Base.

CDP 7.1.7 SP2 and CDP 7.1.9 Components with REST API differences

The following is a list of all the APIs that have been modified or removed in this release of CDP Private Cloud Base. The APIs listed here are on a component by component basis and may include recommendations for how to manage the change.

REST API Compatibility changes for Ranger

Breaking changes to APIs in CDP for Ranger and recommendations for how to handle them.

updateXResource

added the new path request parameter 'id'

API Method

PUT

API Endpoint

/assets/resources/{id}

Reason for Change

To fail the request if invalid resource id is provided. For details refer: RANGER-3883

Recommendation

Continue using with valid resource id

updateServiceDef

added the new path request parameter 'id'

API Method

PUT

API Endpoint

/plugins/definitions/{id}

Reason for Change

To fail the request if invalid service-def id is provided. For details refer: RANGER-3883

Recommendation

Continue using with valid service-def id

updatePolicy

added the new path request parameter 'id'

API Method

PUT

API Endpoint

/plugins/policies/{id}

Reason for Change

To fail the request if invalid policy id is provided. For details refer: RANGER-3883

Recommendation

Continue using with valid policy id

addUsersAndGroups2

added the new path request parameter 'id'

API Method

PUT

API Endpoint

/roles/roles/{id}/addUsersAndGroups

Reason for Change

To fail the request if invalid role id is provided. For details refer: RANGER-3883

Recommendation

Continue using with valid role id

removeAdminFromUsersAndGroups2

added the new path request parameter 'id'

API Method

PUT

API Endpoint

/roles/roles/{id}/removeAdminFromUsersAndGroups

Reason for Change

To fail the request if invalid role id is provided. For details refer: RANGER-3883

Recommendation

Continue using with valid role id

removeUsersAndGroups2

added the new path request parameter 'id'

API Method

PUT

API Endpoint

/roles/roles/{id}/removeUsersAndGroups

Reason for Change

To fail the request if invalid role id is provided. For details refer: RANGER-3883

Recommendation

Continue using with valid role id

suggestUserFirstName

api path removed without deprecation

API Method

GET

API Endpoint

/users/firstnames

Reason for Change

Removed as it was not having implementation. For details refer: RANGER-3885

Recommendation

Removed and not available for use

updateXGroupPermission

added the new path request parameter 'id'

API Method

PUT

API Endpoint

/xusers/permission/group/{id}

Reason for Change

To fail the request if invalid group id is provided. For details refer: RANGER-3883

Recommendation

Continue using with valid group id

CDP 7.1.7 SP2 and CDP 7.1.9 Compatible Components

The REST APIs for the components mentioned are 100% backward compatible between CDP Private Cloud Base 7.1.7 SP2 and CDP Private Cloud Base 7.1.9.

- Atlas
- SMM
- Cruise Control
- SRM
- Schema Registry
- Hive

CDP 7.1.8 CHF 12 and CDP 7.1.9 Components with REST API differences

The following is a list of all the APIs that have been modified or removed in this release of CDP Private Cloud Base. The APIs listed here are on a component by component basis and may include recommendations for how to manage the change.

REST API Compatibility changes for Ranger

Breaking changes to APIs in CDP for Ranger and recommendations for how to handle them.

updateXResource

added the new path request parameter 'id'

API Method

PUT

API Endpoint

/assets/resources/{id}

Reason for Change

To fail the request if invalid resource id is provided. For details refer: RANGER-3883

Recommendation

Continue using with valid resource id

updateServiceDef

added the new path request parameter 'id'

API Method

PUT

API Endpoint

/plugins/definitions/{id}

Reason for Change

To fail the request if invalid service-def id is provided. For details refer: RANGER-3883

Recommendation

Continue using with valid service-def id

updatePolicy

added the new path request parameter 'id'

API Method

PUT

API Endpoint

/plugins/policies/{id}

Reason for Change

To fail the request if invalid policy id is provided. For details refer: RANGER-3883

Recommendation

Continue using with valid policy id

addUsersAndGroups2

added the new path request parameter 'id'

API Method

PUT

API Endpoint

/roles/roles/{id}/addUsersAndGroups

Reason for Change

To fail the request if invalid role id is provided. For details refer: RANGER-3883

Recommendation

Continue using with valid role id

removeAdminFromUsersAndGroups2

added the new path request parameter 'id'

API Method

PUT

API Endpoint

/roles/roles/{id}/removeAdminFromUsersAndGroups

Reason for Change

To fail the request if invalid role id is provided. For details refer: RANGER-3883

Recommendation

Continue using with valid role id

removeUsersAndGroups2

added the new path request parameter 'id'

API Method

PUT

API Endpoint

/roles/roles/{id}/removeUsersAndGroups

Reason for Change

To fail the request if invalid role id is provided. For details refer: RANGER-3883

Recommendation

Continue using with valid role id

suggestUserFirstName

api path removed without deprecation

API Method

GET

API Endpoint

/users/firstnames

Reason for Change

Removed as it was not having implementation. For details refer: RANGER-3885

Recommendation

Removed and not available for use

updateXGroupPermission

added the new path request parameter 'id'

API Method

PUT

API Endpoint

/xusers/permission/group/{id}

Reason for Change

To fail the request if invalid group id is provided. For details refer: RANGER-3883

Recommendation

Continue using with valid group id

Streams Messaging Manager

Breaking changes to APIs in CDP for SMM and recommendations for how to handle them.

For information on the changes to SMM API, see [Deprecation notices in Streams Messaging Manager](#).

CDP 7.1.8 CHF 12 and CDP 7.1.9 Compatible Components

The REST APIs for the components mentioned are 100% backward compatible between CDP Private Cloud Base 7.1.8 and CDP Private Cloud Base 7.1.9.

- Atlas
- Cruise Control
- SRM
- Schema Registry
- Hive

Deprecation notices in Cloudera Runtime 7.1.9

Certain features and functionalities have been removed or deprecated in Cloudera Runtime 7.1.9. You must review these items to understand whether you must modify your existing configuration. You can also learn about the features that will be removed or deprecated in the future release to plan for the required changes.

Terminology

Items in this section are designated as follows:

Deprecated

Technology that Cloudera is removing in a future CDP release. Marking an item as deprecated gives you time to plan for removal in a future CDP release.

Moving

Technology that Cloudera is moving from a future CDP release and is making available through an alternative Cloudera offering or subscription. Marking an item as moving gives you time to plan for removal in a future CDP release and plan for the alternative Cloudera offering or subscription for the technology.

Removed

Technology that Cloudera has removed from CDP and is no longer available or supported as of this release. Take note of technology marked as removed since it can potentially affect your upgrade plans.

You must contact Cloudera Support or your Cloudera Account Team if you have any questions.

Platform and OS

The listed Operating Systems, databases, and instant client library are deprecated or removed from the 7.1.9 release.

Database Support:

The listed databases are deprecated from the 7.1.9 release.

- Postgres 10
- MariaDB 10.3
- MariaDB 10.2
- Oracle 12
- MySQL 5.6

Operating System

The listed operating system is removed from the 7.1.9 release.

- Ubuntu 18.04



Note: Cloudera Manager 7.11.3 does not support the Ubuntu 18 Operating System. You must use the following path before you upgrade to CDP Private Cloud Base 7.1.9:

1. Upgrade the Operating System from Ubuntu 18 to Ubuntu 20. For performing major OS upgrade, see [Upgrading the Operating System to a new Major Version](#).
2. Upgrade Cloudera Manager to Cloudera Manager 7.11.3 CHF4.
3. Upgrade the CDP cluster to CDP Private Cloud Base 7.1.9.

Removed Oracle Instant Client library from Cloudera Archive

The Oracle Instant Client library is no longer available on the Cloudera Archive for download. You must download the Oracle Instant Client library (both basic and SDK clients) from the [Oracle website](#). Oracle Instant Client library is needed if you want to use Oracle as a backend database for your Runtime components such as Hue.

Removed Support for IBM PPC OS parcels

Support for IBM PPC OS parcels is no longer available from version 7.1.9 CHF 3 onwards.

Deprecation Notices for Spark 2

Spark 2 is deprecated in Cloudera Runtime 7.1.9, therefore 7.1.9 is the last runtime release where Spark 2 is supported. You will need to migrate your Spark 2 applications to Spark 3.3.2. You must ensure that your jobs are Spark 3.3.2 compliant. Please contact Cloudera Support or your Cloudera Account Team if you have any questions.

Deprecated

Spark 2

Since 7.1.9 is the last version in which Spark 2 is supported, you will need to migrate to Spark 3.3.2 before you upgrade to a later version. See [Updating Spark 2 applications for Spark 3](#) linked below.



Note:

Using Spark 2 with Ubuntu 22.04 (supported from CDP Private Cloud version 7.1.9 SP1) **requires installing Python 3.7**.

The highest Python version supported with Spark 2 is 3.7, but Ubuntu 22.04 comes with Python 3.10 installed as default.

Related Information

[Using Spark 2 applications for Spark 3](#)

Deprecation Notice for DAS

Data Analytics Studio (DAS) has been deprecated and is no longer available in CDP Private Cloud Base starting with 7.1.9.

Removed component

Hue now replaces DAS. DAS features to support Hive and Tez such as running queries, defining HPL/SQL, the Job Browser, query explorer, query compare, and more, have been migrated to Hue, and the Hue Query Processor. After you upgrade to this release, you will not see the option to add the DAS service to your cluster. Cloudera recommends you use Hue for all use cases where you might have previously used DAS.

Deprecation Notices for Zeppelin

Starting from Cloudera Runtime 7.1.9, Zeppelin has been deprecated. While Zeppelin will continue to receive full support in 7.1.9, we recommend you to consider migrating to other Cloudera products. Detailed guidance on the recommended migration process will be provided shortly. Reach out to Cloudera Support or your dedicated Cloudera Account Team if you have inquiries or require further assistance.

Deprecation Notices for Apache Oozie

Certain features and functionality in Oozie are deprecated or removed in Cloudera Runtime 7.1.9.

Deprecated

Oozie's Spark action

Due to the discontinuation and deprecation of Spark 2 in CDP 7.1.9, Cloudera decided to deprecate Oozie Spark actions, which are based on Spark 2. Consequently, Oozie's Spark actions **will be disabled by default**, and if you attempt to execute a Spark action, an error will be raised.

Starting from 7.1.9, Oozie introduces the new Spark 3 based Spark 3 actions.

Deprecation Notices for Cruise Control

Certain features and functionality in Cruise Control are deprecated or removed in Cloudera Runtime 7.1.9. You must review these changes along with the information about the features in Cruise Control that will be removed or deprecated in a future release.

Deprecated

Cloudera Manager Metrics Reporter

The Cloudera Manager Metrics Reporter is deprecated from Cruise Control and is replaced with the Cruise Control Metrics Reporter.

Deprecation Notices for Apache Kafka

Certain features and functionality in Kafka are deprecated or removed in Cloudera Runtime 7.1.9. You must review these changes along with the information about the features in Kafka that will be removed or deprecated in a future release.



Important: The following list of deprecated and removed items is not exhaustive and only contains items that have a direct and immediate effect on Kafka in CDP. For a full list of deprecation and/or removals in the version Apache Kafka shipped with Runtime, review the *Notable Changes* as well as the *Release Notes* on <https://kafka.apache.org/>.

Removed

Kafka Connect Prometheus Metrics Port (`connect.prometheus.metrics.port`)

This property is removed and replaced by Jetty Metrics Port and Secure Jetty Metrics Port. As a result of this change, the recommended way of setting up Prometheus for Streams Messaging Manager is changed. For more information, see [Setting up Prometheus for Streams Messaging Manager](#).

Deprecated

MirrorMaker (MM1)

MirrorMaker is deprecated. Cloudera recommends that you use Streams Replication Manager (SRM) instead.


--zookeeper

The --zookeeper option is only supported for the kafka-configs tool and should be only used when updating SCRAM Credential configurations. The --zookeeper option is either deprecated in or removed from other Kafka command line tools. Cloudera recommends that you use the --bootstrap-server option instead.

Navigator Key Trustee Server

Navigator Key Trustee Server (KTS) is being deprecated in CDP Private Cloud Base 7.1.9 immediately. Navigator Key Trustee Server will continue to be supported for the entire lifespan of CDP Private Cloud Base 7.1.9, including all service packs. For identical key and encryption management, customers must move to Ranger KMS. Also, Customers using Navigator Encrypt (NavEncrypt) will need to migrate metadata storage to Ranger KMS.

The ability to migrate keys and encryption metadata from Navigator KTS to Ranger KMS exists as of CDP Private Cloud Base 7.1.9. However, some more advanced configurations will require a more recent version of CDP 7.1.9. Note that the version mentioned is of the CDP cluster that Ranger KMS will be installed on and not the version of the KTS cluster; which can remain on a version no older than CDP 7.1.7. Customers using Navigator Key Trustee Server without an HSM, with or without Navigator Encrypt, can migrate to Ranger KMS in CDP 7.1.9. If an HSM is being used, but Navigator Encrypt is not being used, then the customer must be on at least CDP 7.1.9 Service Pack 1 (SP1). If an HSM and Navigator Encrypt is being used, please notify your account team. Refer to the flowchart below for a summary of the minimum required CDP version requirements.



Detailed migration steps are available in the CDP documentation topic [Migrating keys from Key Trustee Server to Ranger KMS](#).

Fixed Common Vulnerabilities and Exposures 7.1.9

Learn more about the Common vulnerabilities and Exposures (CVEs) that were fixed in this release.

- [CVE-2022-31160](#) - jQuery UI is a curated set of user interface interactions, effects, widgets, and themes built on top of jQuery.
- [CVE-2022-23302](#) - JMSSink in all versions of Log4j 1.x is vulnerable to deserialization of untrusted data when the attacker has write access to the Log4j configuration or if the configuration references an LDAP service the attacker has access to.
- [CVE-2022-23307](#) - CVE-2020-9493 identified a deserialization issue that was present in Apache Chainsaw.
- [CVE-2021-4048](#) - An out-of-bounds read flaw was found in the CLARRV, DLARRV, SLARRV, and ZLARRV functions in lapack through version 3.10.0, as also used in OpenBLAS before version 0.3.18.
- [CVE-2020-28462](#) - This affects all versions of package ion-parser. If an attacker submits a malicious INI file to an application that parses it with parse , they will pollute the prototype on the application.
- [CVE-2022-36364](#) - Apache Calcite Avatica JDBC driver creates HTTP client instances based on class names provided via 'httpclient_impl' connection property; however, the driver does not verify if the class implements the expected interface before instantiating it, which can lead to code execution loaded via arbitrary classes and in rare cases remote code execution.

- [CVE-2021-31684](#) - A vulnerability was discovered in the `indexOf` function of `JSONParserByteArray` in JSON Smart versions 1.3 and 2.4 which causes a denial of service (DOS) via a crafted web request.
- [CVE-2022-30187](#) - Azure Storage Library Information Disclosure Vulnerability
- [CVE-2020-26939](#) - In Legion of the Bouncy Castle BC before 1.61 and BC-FJA before 1.0.1.2, attackers can obtain sensitive information about a private exponent because of Observable Differences in Behavior to Error Inputs. This occurs in `org.bouncycastle.crypto.encodings.OAEPEncoding`.
- [CVE-2020-15522](#) - Bouncy Castle BC Java before 1.66, BC C# .NET before 1.8.7, BC-FJA before 1.0.1.2, 1.0.2.1, and BC-FNA before 1.0.1.1 have a timing issue within the EC math library that can expose information about the private key when an attacker is able to observe timing information for the generation of multiple deterministic ECDSA signatures.
- [CVE-2020-0187](#) - In `engineSetMode` of `BaseBlockCipher.java`, there is a possible incorrect cryptographic algorithm chosen due to an incomplete comparison.
- [CVE-2021-29243](#) - Cloudera Manager 5.x, 6.x, 7.1.x, 7.2.x, and 7.3.x allows XSS.
- [CVE-2021-32482](#) - Cloudera Manager 5.x, 6.x, 7.1.x, 7.2.x, and 7.3.x allows XSS via the path parameter.
- [CVE-2021-35515](#) - When reading a specially crafted 7Z archive, the construction of the list of codecs that decompress an entry can result in an infinite loop.
- [CVE-2021-35516](#) - When reading a specially crafted 7Z archive, Compress can be made to allocate large amounts of memory that finally leads to an out of memory error even for very small inputs.
- [CVE-2021-35517](#) - When reading a specially crafted TAR archive, Compress can be made to allocate large amounts of memory that finally leads to an out of memory error even for very small inputs.
- [CVE-2021-36090](#) - When reading a specially crafted ZIP archive, Compress can be made to allocate large amounts of memory that finally leads to an out of memory error even for very small inputs.
- [CVE-2018-11771](#) - When reading a specially crafted ZIP archive, the read method of Apache Commons Compress 1.7 to 1.17's `ZipArchiveInputStream` can fail to return the correct EOF indication after the end of the stream has been reached.
- [CVE-2012-5783](#) - Apache Commons HttpClient 3.x, as used in Amazon Flexible Payments Service (FPS) merchant Java SDK and other products, does not verify that the server hostname matches a domain name in the subject's Common Name (CN) or subjectAltName field of the X.509 certificate, which allows man-in-the-middle attackers to spoof SSL servers via an arbitrary valid certificate.
- [CVE-2021-42550](#) - In logback version 1.2.7 and prior versions, an attacker with the required privileges to edit configurations files could craft a malicious configuration allowing to execute arbitrary code loaded from LDAP servers.
- [CVE-2021-26291](#) - Apache Maven will follow repositories that are defined in a dependency's Project Object Model (pom) which may be surprising to some users, resulting in potential risk if a malicious actor takes over that repository or is able to insert themselves into a position to pretend to be that repository.
- [CVE-2021-28170](#) - In the Jakarta Expression Language implementation 3.0.3 and earlier, a bug in the `ELParserTokenManager` enables invalid EL expressions to be evaluated as if they were valid.
- [CVE-2019-19919](#) - Versions of handlebars prior to 4.3.0 are vulnerable to Prototype Pollution leading to Remote Code Execution. Templates may alter an Object's `__proto__` and `__defineGetter__` properties, which may allow an attacker to execute arbitrary code through crafted payloads.
- [CVE-2022-25647](#) - The package `com.google.code.gson:gson` before 2.8.9 are vulnerable to Deserialization of Untrusted Data via the `writeReplace()` method in internal classes, which may lead to DoS attacks.
- [CVE-2019-10219](#) - A vulnerability was found in Hibernate-Validator. The `SafeHtml` validator annotation fails to properly sanitize payloads consisting of potentially malicious code in HTML comments and instructions. This vulnerability can result in an XSS attack.
- [CVE-2020-10693](#) - A flaw was found in Hibernate Validator version 6.1.2.Final. A bug in the message interpolation processor enables invalid EL expressions to be evaluated as if they were valid.
- [CVE-2018-1000840](#) - Processing Foundation Processing version 3.4 and earlier contains a XML External Entity (XXE) vulnerability in `loadXML()` function that can result in An attacker can read arbitrary files and exfiltrate their contents via HTTP requests.
- [CVE-2021-42357](#) - When using Apache Knox SSO prior to 1.6.1, a request could be crafted to redirect a user to a malicious page due to improper URL parsing.

- [CVE-2017-1000028](#) - Oracle, GlassFish Server Open Source Edition 4.1 is vulnerable to both authenticated and unauthenticated Directory Traversal vulnerability, that can be exploited by issuing a specially crafted HTTP GET request.
- [CVE-2014-0075](#) - Integer overflow in the parseChunkHeader function in java/org/apache/coyote/http11/filters/ChunkedInputFilter.java in Apache Tomcat before 6.0.40, 7.x before 7.0.53, and 8.x before 8.0.4 allows remote attackers to cause a denial of service (resource consumption) via a malformed chunk size in chunked transfer coding of a request during the streaming of data.
- [CVE-2014-0099](#) - Integer overflow in java/org/apache/tomcat/util/buf/Ascii.java in Apache Tomcat before 6.0.40, 7.x before 7.0.53, and 8.x before 8.0.4, when operated behind a reverse proxy, allows remote attackers to conduct HTTP request smuggling attacks via a crafted Content-Length HTTP header.
- [CVE-2015-6584](#) - Cross-site scripting (XSS) vulnerability in the DataTables plugin 1.10.8 and earlier for jQuery allows remote attackers to inject arbitrary web script or HTML via the scripts parameter to media/unit_testing/templates/6776.php.
- [CVE-2020-13949](#) - In Apache Thrift 0.9.3 to 0.13.0, malicious RPC clients could send short messages which would result in a large memory allocation, potentially leading to denial of service.
- [CVE-2021-44548](#) - An Improper Input Validation vulnerability in DataImportHandler of Apache Solr allows an attacker to provide a Windows UNC path resulting in an SMB network call being made from the Solr host to another host on the network.
- [CVE-2021-34538](#) - Apache Hive before 3.1.3 "CREATE" and "DROP" function operations does not check for necessary authorization of involved entities in the query. It was found that an unauthorized user can manipulate an existing UDF without having the privileges to do so.
- [CVE-2018-18928](#) - International Components for Unicode (ICU) for C/C++ 63.1 has an integer overflow in number::impl::DecimalQuantity::toScientificString() in i18n/number_decimalquantity.cpp.
- [CVE-2020-10531](#) - An issue was discovered in International Components for Unicode (ICU) for C/C++ through 66.1. An integer overflow, leading to a heap-based buffer overflow, exists in the UnicodeString::doAppend() function in common/unistr.cpp.
- [CVE-2020-21913](#) - International Components for Unicode (ICU-20850) v66.1 was discovered to contain a use after free bug in the pkg_createWithAssemblyCode function in the file tools/pkgdata/pkgdata.cpp.
- [CVE-2020-25649](#) - A flaw was found in FasterXML Jackson Databind, where it did not have entity expansion secured properly. This flaw allows vulnerability to XML external entity (XXE) attacks. The highest threat from this vulnerability is data integrity.
- [CVE-2020-28491](#) - This affects the package com.fasterxml.jackson.dataformat:jackson-dataformat-cbor from 0 and before 2.11.4, from 2.12.0-rc1 and before 2.12.1. Unchecked allocation of byte buffer can cause a java.lang.OutOfMemoryError exception.
- [CVE-2021-28168](#) - Eclipse Jersey 2.28 to 2.33 and Eclipse Jersey 3.0.0 to 3.0.1 contains a local information disclosure vulnerability. This is due to the use of the File.createTempFile which creates a file inside of the system temporary directory with the permissions: -rw-r--r--.
- [CVE-2022-36033](#) - jsoup is a Java HTML parser, built for HTML editing, cleaning, scraping, and cross-site scripting (XSS) safety. jsoup may incorrectly sanitize HTML including `javascript:` URL expressions, which could allow XSS attacks when a reader subsequently clicks that link.
- [CVE-2019-17571](#) - included in Log4j 1.2 is a SocketServer class that is vulnerable to deserialization of untrusted data which can be exploited to remotely execute arbitrary code when combined with a deserialization gadget when listening to untrusted network traffic for log data.
- [CVE-2021-4104](#) - JMSAppender in Log4j 1.2 is vulnerable to deserialization of untrusted data when the attacker has write access to the Log4j configuration.
- [CVE-2017-12629](#) - Remote code execution occurs in Apache Solr before 7.1 with Apache Lucene before 7.1 by exploiting XXE in conjunction with use of a Config API add-listener command to reach the RunExecutableListener class. Elasticsearch, although it uses Lucene, is NOT vulnerable to this.
- [CVE-2022-24614](#) - When reading a specially crafted JPEG file, metadata-extractor up to 2.16.0 can be made to allocate large amounts of memory that finally leads to an out-of-memory error even for very small inputs.
- [CVE-2017-18214](#) - The moment module before 2.19.3 for Node.js is prone to a regular expression denial of service via a crafted date string, a different vulnerability than CVE-2016-4055.
- [CVE-2021-20328](#) - Specific versions of the Java driver that support client-side field level encryption (CSFLE) fail to perform correct host name verification on the KMS server's certificate.

- [CVE-2011-1797](#) - WebKit, as used in Apple Safari before 5.0.6, allows remote attackers to execute arbitrary code or cause a denial of service (memory corruption and application crash) via a crafted web site, a different vulnerability than other WebKit CVEs listed in APPLE-SA-2011-07-20-1.
- [CVE-2021-0341](#) - In verifyHostName of OkHostnameVerifier.java, there is a possible way to accept a certificate for the wrong domain due to improperly used crypto. This could lead to remote information disclosure with no additional execution privileges needed.
- [CVE-2021-42575](#) - The OWASP Java HTML Sanitizer before 20211018.1 does not properly enforce policies associated with the SELECT, STYLE, and OPTION elements.
- [CVE-2022-26336](#) - A shortcoming in the HMEF package of poi-scratchpad (Apache POI) allows an attacker to cause an Out of Memory exception.
- [CVE-2022-21724](#) - pgjdbc is the official PostgreSQL JDBC Driver. A security hole was found in the jdbc driver for postgresql database while doing security research.
- [CVE-2022-26520](#) - ** DISPUTED ** In pgjdbc before 42.3.3, an attacker (who controls the jdbc URL or properties) can call java.util.logging.FileHandler to write to arbitrary files through the loggerFile and loggerLevel connection properties.
- [CVE-2022-31197](#) - PostgreSQL JDBC Driver (PgJDBC for short) allows Java programs to connect to a PostgreSQL database using standard, database independent Java code.
- [CVE-2012-6708](#) - jQuery before 1.9.0 is vulnerable to Cross-site Scripting (XSS) attacks.
- [CVE-2015-9251](#) - jQuery before 3.0.0 is vulnerable to Cross-site Scripting (XSS) attacks when a cross-domain Ajax request is performed without the dataType option, causing text/javascript responses to be executed.
- [CVE-2019-11358](#) - jQuery before 3.4.0, as used in Drupal, Backdrop CMS, and other products, mishandles jQuery.extend(true, {}, ...) because of Object.prototype pollution.
- [CVE-2020-11023](#) - In jQuery versions greater than or equal to 1.0.3 and before 3.5.0, passing HTML containing <option> elements from untrusted sources - even after sanitizing it - to one of jQuery's DOM manipulation methods (i.e. .html(), .append(), and others) may execute untrusted code.
- [CVE-2020-7656](#) - jquery prior to 1.9.0 allows Cross-site Scripting attacks via the load method. The load method fails to recognize and remove "<script>" HTML tags that contain a whitespace character, i.e: "</script >", which results in the enclosed script logic to be executed.
- [CVE-2016-7103](#) - Cross-site scripting (XSS) vulnerability in jQuery UI before 1.12.0 might allow remote attackers to inject arbitrary web script or HTML via the closeText parameter of the dialog function.
- [CVE-2021-41182](#) - jQuery-UI is the official jQuery user interface library. Prior to version 1.13.0, accepting the value of the `altField` option of the Datepicker widget from untrusted sources may execute untrusted code.
- [CVE-2021-41183](#) - jQuery-UI is the official jQuery user interface library. Prior to version 1.13.0, accepting the value of various `\*Text` options of the Datepicker widget from untrusted sources may execute untrusted code.
- [CVE-2021-41184](#) - jQuery-UI is the official jQuery user interface library. Prior to version 1.13.0, accepting the value of the `of` option of the .position() util from untrusted sources may execute untrusted code.
- [CVE-2010-5312](#) - Cross-site scripting (XSS) vulnerability in jquery.ui.dialog.js in the Dialog widget in jQuery UI before 1.10.0 allows remote attackers to inject arbitrary web script or HTML via the title option.
- [CVE-2020-11022](#) - In jQuery versions greater than or equal to 1.2 and before 3.5.0, passing HTML from untrusted sources - even after sanitizing it - to one of jQuery's DOM manipulation methods (i.e. .html(), .append(), and others) may execute untrusted code.
- [CVE-2017-15288](#) - The compilation daemon in Scala before 2.10.7, 2.11.x before 2.11.12, and 2.12.x before 2.12.4 uses weak permissions for private files in /tmp/scala-devel/\${USER:shared}/scalac-compile-server-port, which allows local users to write to arbitrary class files and consequently gain privileges.
- [CVE-2022-33891](#) - The Apache Spark UI offers the possibility to enable ACLs via the configuration option spark.acls.enable. With an authentication filter, this checks whether a user has access permissions to view or modify the application.
- [CVE-2021-38296](#) - Apache Spark supports end-to-end encryption of RPC connections via "spark.authenticate" and "spark.network.crypto.enabled". In versions 3.1.2 and earlier, it uses a bespoke mutual authentication protocol that allows for full encryption key recovery.
- [CVE-2018-1270](#) - Spring Framework, versions 5.0 prior to 5.0.5 and versions 4.3 prior to 4.3.15 and older unsupported versions, allow applications to expose STOMP over WebSocket endpoints with a simple, in-memory STOMP broker through the spring-messaging module.

- [CVE-2022-22965](#) - A Spring MVC or Spring WebFlux application running on JDK 9+ may be vulnerable to remote code execution (RCE) via data binding. The specific exploit requires the application to run on Tomcat as a WAR deployment. If the application is deployed as a Spring Boot executable jar, i.e. the default, it is not vulnerable to the exploit. However, the nature of the vulnerability is more general, and there may be other ways to exploit it.
- [CVE-2015-5211](#) - Under some situations, the Spring Framework 4.2.0 to 4.2.1, 4.0.0 to 4.1.7, 3.2.0 to 3.2.14 and older unsupported versions is vulnerable to a Reflected File Download (RFD) attack.
- [CVE-2016-5007](#) - Both Spring Security 3.2.x, 4.0.x, 4.1.0 and the Spring Framework 3.2.x, 4.0.x, 4.1.x, 4.2.x rely on URL pattern mappings for authorization and for mapping requests to controllers respectively. Differences in the strictness of the pattern matching mechanisms, for example with regards to space trimming in path segments, can lead Spring Security to not recognize certain paths as not protected that are in fact mapped to Spring MVC controllers that should be protected.
- [CVE-2016-9878](#) - An issue was discovered in Pivotal Spring Framework before 3.2.18, 4.2.x before 4.2.9, and 4.3.x before 4.3.5. Paths provided to the ResourceServlet were not properly sanitized and as a result exposed to directory traversal attacks.
- [CVE-2018-11040](#) - Spring Framework, versions 5.0.x prior to 5.0.7 and 4.3.x prior to 4.3.18 and older unsupported versions, allows web applications to enable cross-domain requests via JSONP (JSON with Padding) through AbstractJsonpResponseBodyAdvice for REST controllers and MappingJackson2JsonView for browser requests.
- [CVE-2018-1257](#) - Spring Framework, versions 5.0.x prior to 5.0.6, versions 4.3.x prior to 4.3.17, and older unsupported versions allows applications to expose STOMP over WebSocket endpoints with a simple, in-memory STOMP broker through the spring-messaging module. A malicious user (or attacker) can craft a message to the broker that can lead to a regular expression, denial of service attack.
- [CVE-2020-5421](#) - In Spring Framework versions 5.2.0 - 5.2.8, 5.1.0 - 5.1.17, 5.0.0 - 5.0.18, 4.3.0 - 4.3.28, and older unsupported versions, the protections against RFD attacks from CVE-2015-5211 may be bypassed depending on the browser used through the use of a jsessionid path parameter.
- [CVE-2022-22950](#) - In Spring Framework versions 5.3.0 - 5.3.16 and older unsupported versions, it is possible for a user to provide a specially crafted SpEL expression that may cause a denial of service condition.
- [CVE-2018-11039](#) - Spring Framework (versions 5.0.x prior to 5.0.7, versions 4.3.x prior to 4.3.18, and older unsupported versions) allow web applications to change the HTTP request method to any HTTP method (including TRACE) using the HiddenHttpMethodFilter in Spring MVC.
- [CVE-2015-3192](#) - Pivotal Spring Framework before 3.2.14 and 4.x before 4.1.7 do not properly process inline DTD declarations when DTD is not entirely disabled, which allows remote attackers to cause a denial of service (memory consumption and out-of-memory errors) via a crafted XML file.
- [CVE-2022-22968](#) - In Spring Framework versions 5.3.0 - 5.3.18, 5.2.0 - 5.2.20, and older unsupported versions, the patterns for disallowedFields on a DataBinder are case sensitive which means a field is not effectively protected unless it is listed with both upper and lower case for the first character of the field, including upper and lower case for the first character of all nested fields within the property path.
- [CVE-2022-22970](#) - In spring framework versions prior to 5.3.20+ , 5.2.22+ and old unsupported versions, applications that handle file uploads are vulnerable to DoS attack if they rely on data binding to set a MultipartFile or javax.servlet.Part to a field in a model object.
- [CVE-2022-22971](#) - In spring framework versions prior to 5.3.20+ , 5.2.22+ and old unsupported versions, application with a STOMP over WebSocket endpoint is vulnerable to a denial of service attack by an authenticated user.
- [CVE-2022-22978](#) - In spring security versions prior to 5.4.11+ , 5.5.7+ , 5.6.4+ and older unsupported versions, RegexRequestMatcher can easily be misconfigured to be bypassed on some servlet containers. Applications using RegexRequestMatcher with `.` in the regular expression are possibly vulnerable to an authorization bypass.
- [CVE-2021-22112](#) - Spring Security 5.4.x prior to 5.4.4, 5.3.x prior to 5.3.8.RELEASE, 5.2.x prior to 5.2.9.RELEASE, and older unsupported versions can fail to save the SecurityContext if it is changed more than once in a single request. A malicious user cannot cause the bug to happen (it must be programmed in).
- [CVE-2016-9879](#) - An issue was discovered in Pivotal Spring Security before 3.2.10, 4.1.x before 4.1.4, and 4.2.x before 4.2.1. Spring Security does not consider URL path parameters when processing security constraints.
- [CVE-2019-11272](#) - Spring Security, versions 4.2.x up to 4.2.12, and older unsupported versions support plain text passwords using PlaintextPasswordEncoder. If an application using an affected version of Spring Security is

leveraging PlaintextPasswordEncoder and a user has a null encoded password, a malicious user (or attacker) can authenticate using a password of "null".

- [CVE-2019-3795](#) - Spring Security versions 4.2.x prior to 4.2.12, 5.0.x prior to 5.0.12, and 5.1.x prior to 5.1.5 contain an insecure randomness vulnerability when using SecureRandomFactoryBean#setSeed to configure a SecureRandom instance.
- [CVE-2022-22976](#) - Spring Security versions 5.5.x prior to 5.5.7, 5.6.x prior to 5.6.4, and earlier unsupported versions contain an integer overflow vulnerability. When using the BCrypt class with the maximum work factor (31), the encoder does not perform any salt rounds, due to an integer overflow error. The default settings are not affected by this CVE.
- [CVE-2020-5408](#) - Spring Security versions 5.3.x prior to 5.3.2, 5.2.x prior to 5.2.4, 5.1.x prior to 5.1.10, 5.0.x prior to 5.0.16 and 4.2.x prior to 4.2.16 use a fixed null initialization vector with CBC Mode in the implementation of the queryable text encryptor.
- [CVE-2016-1000027](#) - Pivotal Spring Framework through 5.3.16 suffers from a potential remote code execution (RCE) issue if used for Java deserialization of untrusted data.
- [CVE-2022-25169](#) - The BPG parser in versions of Apache Tika before 1.28.2 and 2.4.0 may allocate an unreasonable amount of memory on carefully crafted files.
- [CVE-2022-30126](#) - In Apache Tika, a regular expression in our StandardsText class, used by the StandardsExtractingContentHandler could lead to a denial of service caused by backtracking on a specially crafted file.
- [CVE-2022-33879](#) - The initial fixes in CVE-2022-30126 and CVE-2022-30973 for regexes in the StandardsExtractingContentHandler were insufficient, and we found a separate, new regex DoS in a different regex in the StandardsExtractingContentHandler.
- [CVE-2022-34305](#) - In Apache Tomcat 10.1.0-M1 to 10.1.0-M16, 10.0.0-M1 to 10.0.22, 9.0.30 to 9.0.64 and 8.5.50 to 8.5.81 the Form authentication example in the examples web application displayed user provided data without filtering, exposing a XSS vulnerability.
- [CVE-2016-10735](#) - In Bootstrap 3.x before 3.4.0 and 4.x-beta before 4.0.0-beta.2, XSS is possible in the data-target attribute, a different vulnerability than CVE-2018-14041.
- [CVE-2018-14040](#) - In Bootstrap before 4.1.2, XSS is possible in the collapse data-parent attribute.
- [CVE-2018-14041](#) - In Bootstrap before 4.1.2, XSS is possible in the data-target property of scrollspy.
- [CVE-2018-14042](#) - In Bootstrap before 4.1.2, XSS is possible in the data-container property of tooltip.
- [CVE-2019-8331](#) - In Bootstrap before 3.4.1 and 4.3.x before 4.3.1, XSS is possible in the tooltip or popover data-template attribute.
- [CVE-2021-3642](#) - A flaw was found in Wildfly Elytron in versions prior to 1.10.14.Final, prior to 1.15.5.Final and prior to 1.16.1.Final where ScramServer may be susceptible to Timing Attack if enabled.
- [CVE-2022-23437](#) - There's a vulnerability within the Apache Xerces Java (XercesJ) XML parser when handling specially crafted XML document payloads.
- [CVE-2017-10355](#) - Vulnerability in the Java SE, Java SE Embedded, JRockit component of Oracle Java SE (subcomponent: Networking).
- [CVE-2019-10095](#) - bash command injection vulnerability in Apache Zeppelin allows an attacker to inject system commands into Spark interpreter settings.
- [CVE-2020-13929](#) - Authentication bypass vulnerability in Apache Zeppelin allows an attacker to bypass Zeppelin authentication mechanism to act as another user. This issue affects Apache Zeppelin Apache Zeppelin version 0.9.0 and prior versions.
- [CVE-2021-27578](#) - Cross Site Scripting vulnerability in markdown interpreter of Apache Zeppelin allows an attacker to inject malicious scripts. This issue affects Apache Zeppelin Apache Zeppelin versions prior to 0.9.0.
- [CVE-2021-43138](#) - In Async before 2.6.4 and 3.x before 3.2.2, a malicious user can obtain privileges via the mapValues() method, aka lib/internal/iterator.js createObjectIterator prototype pollution.
- [CVE-2020-28469](#) - This affects the package glob-parent before 5.1.2. The enclosure regex used to check for strings ending in enclosure containing path separator.
- [CVE-2020-7598](#) - minimist before 1.2.2 could be tricked into adding or modifying properties of Object.prototype using a "constructor" or "__proto__" payload.
- [CVE-2022-21803](#) - This affects the package nconf before 0.11.4. When using the memory engine, it is possible to store a nested JSON representation of the configuration.

Important CVE fixes

- [CVE-2022-36944](#) - Scala 2.13.x before 2.13.9 has a Java deserialization chain in its JAR file. On its own, it cannot be exploited.
- [CVE-2022-41881](#) - Netty project is an event-driven asynchronous network application framework. In versions prior to 4.1.86.Final, a `StackOverflowError` can be raised when parsing a malformed crafted message due to an infinite recursion.
- [CVE-2022-41915](#) - Netty project is an event-driven asynchronous network application framework. Starting in version 4.1.83.Final and prior to 4.1.86.Final, when calling `DefaultHttpHeaders.set`` with an `_iterator_` of values, header value validation was not performed, allowing malicious header values in the iterator to perform HTTP Response Splitting.
- [CVE-2021-3642](#) - A flaw was found in Wildfly Elytron in versions prior to 1.10.14.Final, prior to 1.15.5.Final and prior to 1.16.1.
- [CVE-2022-3143](#) - wildfly-elytron: possible timing attacks via use of unsafe comparator. A flaw was found in Wildfly-elytron. Wildfly-elytron uses `java.util.Arrays.equals` in several places, which is unsafe and vulnerable to timing attacks.
- [CVE-2023-20861](#) - In Spring Framework versions 6.0.0 - 6.0.6, 5.3.0 - 5.3.25, 5.2.0.RELEASE - 5.2.22.RELEASE, and older unsupported versions, it is possible for a user to provide a specially crafted SpEL expression that may cause a denial-of-service (DoS) condition.
- [CVE-2023-20860](#) - Spring Framework running version 6.0.0 - 6.0.6 or 5.3.0 - 5.3.25 using `***` as a pattern in Spring Security configuration with the `mvcRequestMatcher` creates a mismatch in pattern matching between Spring Security and Spring MVC, and the potential for a security bypass.
- [CVE-2023-20863](#) - In spring framework versions prior to 5.2.24 release+ ,5.3.27+ and 6.0.8+ , it is possible for a user to provide a specially crafted SpEL expression that may cause a denial-of-service (DoS) condition.
- [CVE-2022-42252](#) - If Apache Tomcat 8.5.0 to 8.5.82, 9.0.0-M1 to 9.0.67, 10.0.0-M1 to 10.0.26 or 10.1.0-M1 to 10.1.0 was configured to ignore invalid HTTP headers via setting `rejectIllegalHeader` to false (the default for 8.5.x only), Tomcat did not reject a request containing an invalid Content-Length header making a request smuggling attack possible if Tomcat was located behind a reverse proxy that also failed to reject the request with the invalid header.
- [CVE-2022-34305](#) - In Apache Tomcat 10.1.0-M1 to 10.1.0-M16, 10.0.0-M1 to 10.0.22, 9.0.30 to 9.0.64 and 8.5.50 to 8.5.81 the Form authentication example in the examples web application displayed user provided data without filtering, exposing a XSS vulnerability.
- [CVE-2022-45143](#) - The `JsonErrorReportValve` in Apache Tomcat 8.5.83, 9.0.40 to 9.0.68 and 10.1.0-M1 to 10.1.1 did not escape the type, message or description values. In some circumstances these are constructed from user provided data and it was therefore possible for users to supply values that invalidated or manipulated the JSON output.
- [CVE-2021-41182](#) - Query-UI is the official jQuery user interface library. Prior to version 1.13.0, accepting the value of the ``altField`` option of the Datepicker widget from untrusted sources may execute untrusted code.
- [CVE-2021-41183](#) - jQuery-UI is the official jQuery user interface library. Prior to version 1.13.0, accepting the value of various ``*Text`` options of the Datepicker widget from untrusted sources may execute untrusted code.
- [CVE-2021-41184](#) - jQuery-UI is the official jQuery user interface library. Prior to version 1.13.0, accepting the value of the ``of`` option of the `.position()` util from untrusted sources may execute untrusted code.
- [CVE-2010-5312](#) - Cross-site scripting (XSS) vulnerability in `jquery.ui.dialog.js` in the Dialog widget in jQuery UI before 1.10.0 allows remote attackers to inject arbitrary web script or HTML via the title option.
- [CVE-2016-7103](#) - Cross-site scripting (XSS) vulnerability in jQuery UI before 1.12.0 might allow remote attackers to inject arbitrary web script or HTML via the `closeText` parameter of the dialog function.
- [CVE-2022-41946](#) - pgjdbc is an open source postgresql JDBC Driver. In affected versions a prepared statement using either ``PreparedStatement.setText(int, InputStream)`` or ``PreparedStatement.setBytea(int, InputStream)`` will create a temporary file if the `InputStream` is larger than 2k.
- [CVE-2021-41561](#) - Improper Input Validation vulnerability in Parquet-MR of Apache Parquet allows an attacker to DoS by malicious Parquet files. This issue affects Apache Parquet-MR version 1.9.0 and later versions.
- [CVE-2017-1000028](#) - Oracle, GlassFish Server Open Source Edition 4.1 is vulnerable to both authenticated and unauthenticated Directory Traversal vulnerability, that can be exploited by issuing a specially crafted HTTP GET request.

- [CVE-2020-21913](#) - International Components for Unicode (ICU-20850) v66.1 was discovered to contain a use after free bug in the `pkg_createWithAssemblyCode` function in the file `tools/pkgdata/pkgdata.cpp`.
- [CVE-2021-3711](#) - In order to decrypt SM2 encrypted data an application is expected to call the API function `EVP_PKEY_decrypt()`. Typically an application will call this function twice.
- [CVE-2021-3712](#) - ASN.1 strings are represented internally within OpenSSL as an `ASN1_STRING` structure which contains a buffer holding the string data and a field holding the buffer length.
- [CVE-2021-44531](#) - Accepting arbitrary Subject Alternative Name (SAN) types, unless a PKI is specifically defined to use a particular SAN type, can result in bypassing name-constrained intermediates.
- [CVE-2022-21824](#) - Due to the formatting logic of the `"console.table()"` function it was not safe to allow user controlled input to be passed to the `"properties"` parameter while simultaneously passing a plain object with at least one property as the first parameter, which could be `"__proto__"`.
- [CVE-2020-13949](#) - In Apache Thrift 0.9.3 to 0.13.0, malicious RPC clients could send short messages which would result in a large memory allocation, potentially leading to denial of service.
- [CVE-2021-23926](#) - The XML parsers used by XMLBeans up to version 2.6.0 did not set the properties needed to protect the user from malicious XML input. Vulnerabilities include possibilities for XML Entity Expansion attacks. Affects XMLBeans up to and including v2.6.0.
- [CVE-2022-23437](#) - There's a vulnerability within the Apache Xerces Java (XercesJ) XML parser when handling specially crafted XML document payloads. This causes, the XercesJ XML parser to wait in an infinite loop, which may sometimes consume system resources for prolonged duration. This vulnerability is present within XercesJ version 2.12.1 and the previous versions.
- [CVE-2023-20861](#) - In Spring Framework versions 6.0.0 - 6.0.6, 5.3.0 - 5.3.25, 5.2.0.RELEASE - 5.2.22.RELEASE, and older unsupported versions, it is possible for a user to provide a specially crafted SpEL expression that may cause a denial-of-service (DoS) condition.
- [CVE-2023-20860](#) - Spring Framework running version 6.0.0 - 6.0.6 or 5.3.0 - 5.3.25 using `"**"` as a pattern in Spring Security configuration with the `mvcRequestMatcher` creates a mismatch in pattern matching between Spring Security and Spring MVC, and the potential for a security bypass.
- [CVE-2023-20863](#) - In spring framework versions prior to 5.2.24 release+ ,5.3.27+ and 6.0.8+ , it is possible for a user to provide a specially crafted SpEL expression that may cause a denial-of-service (DoS) condition.
- [CVE-2022-36364](#) - Apache Calcite Avatica JDBC driver creates HTTP client instances based on class names provided via `'httpClient_impl'` connection property; however, the driver does not verify if the class implements the expected interface before instantiating it, which can lead to code execution loaded via arbitrary classes and in rare cases remote code execution.
- [CVE-2020-25644](#) - A memory leak flaw was found in WildFly OpenSSL in versions prior to 1.1.3.Final, where it removes an HTTP session. It may allow the attacker to cause OOM leading to a denial of service. The highest threat from this vulnerability is to system availability.
- [CVE-2023-28709](#) - The fix for CVE-2023-24998 was incomplete for Apache Tomcat 11.0.0-M2 to 11.0.0-M4, 10.1.5 to 10.1.7, 9.0.71 to 9.0.73 and 8.5.85 to 8.5.87.
- [CVE-2022-45688](#) - A stack overflow in the `XML.toJSONObject` component of `hutool-json` v5.8.10 allows attackers to cause a Denial of Service (DoS) via crafted JSON or XML data.
- [CVE-2023-32697](#) - SQLite JDBC is a library for accessing and creating SQLite database files in Java. `Sqlite-jdbc` addresses a remote code execution vulnerability via JDBC URL. This issue impacting versions 3.6.14.1 through 3.41.2.1 and has been fixed in version 3.41.2.2.
- [CVE-2023-25577](#) - Werkzeug is a comprehensive WSGI web application library. Prior to version 2.2.3, Werkzeug's multipart form data parser will parse an unlimited number of parts, including file parts. Parts can be a small amount of bytes, but each requires CPU time to parse and may use more memory as Python data.
- [CVE-2023-25613](#) - An LDAP Injection vulnerability exists in the `LdapIdentityBackend` of Apache Kerby before 2.0.3.
- [CVE-2022-40153](#) - ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. Reason: This CVE has been rejected as it was incorrectly assigned. All references and descriptions in this candidate have been removed to prevent accidental usage.
- [CVE-2022-40154](#) - ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. Reason: This CVE has been rejected as it was incorrectly assigned. All references and descriptions in this candidate have been removed to prevent accidental usage.

- [CVE-2022-40155](#) - ** REJECT ** DO NOT USE THIS CANDIDATE NUMBER. Reason: This CVE has been rejected as it was incorrectly assigned. All references and descriptions in this candidate have been removed to prevent accidental usage.
- [CVE-2023-22946](#) - In Apache Spark versions prior to 3.4.0, applications using spark-submit can specify a 'proxy-user' to run as, limiting privileges.
- [CVE-2022-25647](#) - The package com.google.code.gson:gson before 2.8.9 are vulnerable to Deserialization of Untrusted Data via the writeReplace() method in internal classes, which may lead to DoS attacks.
- [CVE-2023-22602](#) - When using Apache Shiro before 1.11.0 together with Spring Boot 2.6+, a specially crafted HTTP request may cause an authentication bypass.

Cumulative hotfixes

You can review the list of cumulative hotfixes that were shipped for CDP Private Cloud Base version 7.1.9.

Cumulative hotfix 1

You can review the list of cumulative hotfixes that were shipped for CDP Private Cloud Base version 7.1.9 CHF1.

Cloudera Runtime 7.1.9.2 (Cumulative Hotfix 1) download URL:

Parcel Repository Location

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.2/parcels/
```

Fixed issues in 7.1.9 CHF 1

Know more about the cumulative hotfixes 1 for 7.1.9. This cumulative hotfix was released on November 02, 2023.



Note: Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p2.46689620

- COMPX-15363: Scheduling rules are not restored from CM properly when cluster is restarted
- COMPX-15282: Backport MAPREDUCE-7456 (Extend add-opens flag to container launch commands on JDK17 nodes for YARN)
- COMPX-15236: QM create versions call, creates multiple config_sets
- COMPX-15216: QM Database migration breaks for clusters where H2 database prior to migration had duplicate config_sets for a namespace, version pair
- CDPD-62299: [UnitTest] Oozie unit test failures due to unable to start Hive Metastore server
- CDPD-62232: Hive: Upgrade snappy-java version to 1.1.10.5 in 7.1.9.x
- CDPD-62128: Using centralised version of snappy-java in Search
- CDPD-62126: Using centralised version of snappy-java in Solr
- CDPD-62125: Kafka - Upgrade snappy-java to 1.1.10.5 due to CVE-2023-43642
- CDPD-61810: Datanucleus upgrade causes test failures in Oozie
- CDPD-61672: Fetch JAVA_OPTS variable in db_setup.py which is provided from CSD.
- CDPD-61616: Extend Java opts for Livy to support JDK17 + Isilon
- CDPD-61605: Extend Java opts for Spark to support JDK17 + Isilon
- CDPD-61595: Backport HIVE-27213 to CDH-7.1.9.x
- CDPD-61586: Implement a workaround for Cruise Control CPU metric collection failure 7.1.9+
- CDPD-61568: [AUTOSYNC] Snapshot should use snapshot's keyManager in optimizeDirDeletesAndSubmitRequest
- CDPD-61564: Caused by: java.lang.NoClassDefFoundError: org/datanucleus/store/query/cache/QueryCompilationCache

- CDPD-61562: Exclude reload4j library from ranger
- CDPD-61547: Snoop should not close the 'System.out' and 'System.err'
- CDPD-61540: CDH-7.1.9.1-17 shows compiler errors for zeppelin on all the OSES
- CDPD-61525: Excluding groovy from gateway-cloud-binding
- CDPD-61501: "Sync source" filter in User/Group search in Oracle DB used clusters leads to an error
- CDPD-61439: [7.1.9 CHF 1] In Tag-based policy from Ranger Admin UI, Allow Conditions permissions item is not showing services permissions which have enableDenyAndExceptionsInPolicies flag false
- CDPD-61433: [7.1.x]- Ranger CSV Report extract may fail with Null pointer exception
- CDPD-61432: Bump jackson-mapper-asl to 1.9.13-cloudera.4 version
- CDPD-61398: [AUTOSYNC] LegacyReplicationManager: Delete excess unhealthy with force=true
- CDPD-61379: [AUTOSYNC] Avoid overriding finalize() in CodecBuffer
- CDPD-61324: Backport HIVE-25918 to CDH-7.1.9.x
- CDPD-61317: Backport CDPD-61098 to 719 CHF
- CDPD-61316: Backport CDPD-61018 to 718 CHF and 719 CHF
- CDPD-61314: Backport PHOENIX-6560 Rewrite dynamic SQL queries to use PreparedStatement
- CDPD-61310: Backport PHOENIX-7005 Spark connector tests cannot compile with latest Phoenix
- CDPD-61309: Backport PHOENIX-6899 Query limit not enforced in UncoveredIndexRegionScanner
- CDPD-61306: Backport PHOENIX-6916 Cannot handle ranges where start is a prefix of end for desc columns
- CDPD-61292: Add InterfaceAudience.Public annotations to relevant HBase-MCC classes
- CDPD-61269: Backport PHOENIX-6854 Salted global indexes do not work for queries with uncovered columns
- CDPD-61263: Backport IMPALA-11195 to 7.1.9 CHF
- CDPD-61244: Recon - HeatMap UI doesn't load On Disabling & Enable Recon HeatMap using Feature Flag
- CDPD-61232: Backport IMPALA-10829 to 7.1.9 CHF
- CDPD-61223: Backport HIVE-27303 to CDH-7.1.9.x
- CDPD-61221: Backport SPARK-40617 to 719 CHF
- CDPD-61188: Hue build failed because of latest virtualenv version
- CDPD-61170: [7.1.x] - Improve ExportCSV download time
- CDPD-61148: [AUTOSYNC] LegacyReplicationManager: Unhealthy replicas could block under replication handling
- CDPD-61146: Backport CDPD-57831 to 719 CHF
- CDPD-61125: [719] - Ranger KMS junit tests are failing
- CDPD-61108: [7.1.9 CHF] - RangerJSONAuditWriter creates new log file for writing ranger audits as JSON every time there is an Exception
- CDPD-61103: Backport HIVE-22961 to CDH-7.1.9.x
- CDPD-61072: [AUTOSYNC] Add assertions to BlockOutputStream
- CDPD-61067: [AUTOSYNC] Investigate whether listStatus() is correctly iterating cache
- CDPD-61064: [AUTOSYNC] Fix snapdiff output for key modification
- CDPD-61051: [7.1.9 CHF1 CLONE] - [Intermittent] Active NN not getting latest resource mappings from RMS server
- CDPD-61050: [ranger][replication] empty export roles file causing transform step to fail
- CDPD-61046: Bump NodeJS version to 20.5.1 due to multiple CVEs
- CDPD-61033: Backport HIVE-27632 to CDH-7.1.9.x
- CDPD-61007: Backport HIVE-27304 to CDH-7.1.9.x
- CDPD-61001: Backport HIVE-25576 to 7.1.9.x
- CDPD-60984: [719 CHF1] Ranger - Upgrade Tomcat to 8.5.93/9.0.80 due to CVE-2023-41080
- CDPD-60973: livy_unittests failed in livy-server module
- CDPD-60961: Ozone replication manager uses mismatched replicas as replication sources
- CDPD-60960: Ozone replication manager cannot progress when all nodes have a replica
- CDPD-60951: [7.1.9 CHF1] Add server side validation for service audit filter

- CDPD-60919: [7.1.9 CHF1] [Ranger React UI] Difference in user lookup API request in permissions module page between React UI and BackBone UI
- CDPD-60915: [7.1.9 CHF1] Update swagger version in Ranger
- CDPD-60911: Knox Readiness Awareness and Notification
- CDPD-60876: Ranger Junit Tests failing
- CDPD-60871: [UnitTest] testQueueSizeAfterNormalSubmission fails with 'Too few elements in the queue'
- CDPD-60859: Enable nashorn features in GraalVM
- CDPD-60847: Kafka_connect_ext - Vulnerable Guava version coming from debezium-core:1.9.7.Final
- CDPD-60842: [7.1.9 CHF1] - Fix to use "public/v2/api/zone-headers" api to get list of zones in Access Logs and Report pages
- CDPD-60839: Upgrade Groovy version >= 3.0.8 to support KnoxShell on JDK17 cluster
- CDPD-60817: IMPALA-12409 Don't allow EXTERNAL Iceberg tables to point another Iceberg table in Hive catalog
- CDPD-60794: [7.1.9 CHF1] In Audit, Plugin Status tab if the record of respective service is in second page then Service Type filter for that service would show no result
- CDPD-60772: IMPALA-10086 SqlCastException when comparing char with varchar
- CDPD-60767: [AUTOSYNC] Snapshot Chain corruption because snapshot chain need not be created in increasing order of CreatedTime
- CDPD-60760: [IBM-PPC] hive server2 service is going down after restart on RHEL8.6
- CDPD-60733: KC qe tests should configure Ranger port and protocol in secure and unsecure clusters differently
- CDPD-60728: [AUTOSYNC] Log EC Replica details if a block cannot be read during reconstruction
- CDPD-60722: Backport HIVE-27586 to 7.1.9.x
- CDPD-60718: Solr Initialisation failing for connection to Solr server while loading heatmap
- CDPD-60687: [7.1.x] Ranger - Upgrade Spring Security to 5.7.10/5.8.5/6.0.5/6.1.2 due to CVE-2023-34034 and CVE-2023-34035
- CDPD-60633: [7.1.9 CHF1] Need to fix zone drop-down option in policy listing for user not having 'Security Zone' module permission
- CDPD-60620: Ozone Recon HeatMap - DrillDown to particular volume not working in multiple volumes starting with same initials are present
- CDPD-60608: Ozone Recon HeatMap - Throws 500 when selecting entity type as Volume
- CDPD-60601: OM restart fails due snapshot chain corruption
- CDPD-60598: Ratis-1868. Handling Netty back pressure when streaming ratis log
- CDPD-60591: [AUTOSYNC] EC: Mark EC containers unhealthy when not missing but unrecoverable
- CDPD-60584: [7.1.9 CHF1 CLONE] - Addressing Vulnerability Type:HTTP Security Header Not Detected only for default HTTPS Port 8484 of RMS
- CDPD-60551: FIPS/FISMA: Oozie needs to grab common JVM settings from hadoop-env.sh by using HADOOP_CLIENT_OPTS
- CDPD-60366: [AUTOSYNC] Native library loader fails when system property "native.lib.tmp.dir" is not set
- CDPD-60363: [AUTOSYNC] A mis replicated EC container with UNHEALTHY replicas may not get resolved
- CDPD-60267: Backport HIVE-27595 to CDP
- CDPD-60240: [AUTOSYNC] Improve debug logging in SCMCommonPlacementPolicy when validating nodes
- CDPD-60199: HMS memory leak because of datanucleus-api-jdo bug
- CDPD-60160: Schema Registry Atlas integration does not work with Oracle DB
- CDPD-60072: [AUTOSYNC] Rename should throw exception upon error
- CDPD-60036: FISMA - Solr is accepting ciphers outside FIPS compliant list on port 8985 - 'TLS_DHE_RSA_WITH_AES_256_GCM_SHA384', 'TLS_DHE_RSA_WITH_AES_256_GCM', 'TLS_DHE_RSA_WITH_AES_128_GCM_SHA384', 'TLS_DHE_RSA_WITH_AES_128_GCM'
- CDPD-60022: Can't set volume space quota on volume if volume has linked bucket
- CDPD-59988: [AUTOSYNC] Decommissioning blocked because of under replicated EC containers
- CDPD-59747: Container stuck in QUASI_CLOSED state causing re-replication failure
- CDPD-59683: Ranger audits are not produced when --all option is added while listing the volume

- CDPD-59623: Cruise Control - Upgrade Okio to 3.4.0 due to CVE-2023-3635
- CDPD-59621: Kafka Connect - Upgrade Okio to 3.4.0 due to CVE-2023-3635
- CDPD-59620: Ranger - Upgrade Okio to 3.4.0 due to CVE-2023-3635
- CDPD-59618: Hadoop - Upgrade Okio to 3.4.0 due to CVE-2023-3635
- CDPD-59614: Backport PHOENIX-6952 Do not disable normalizer on salted tables
- CDPD-59482: [UnitTest] testThereAreNoToManyIdenticalCallbackUrlList fails due to host list size assertion
- CDPD-59480: [UnitTest] testQueueSizeWithDelayedElements Oozie unit test fails intermittently with AssertionFailedError
- CDPD-59421: [719] Knox restart failed due to failure in "wait until.." script for cdp-proxy-api which is due to EOFException
- CDPD-59379: Backport CDPD-58191 to 7.1.x CHFs
- CDPD-59344: Fix and backport PHOENIX-6999 Point lookups fail with reverse scan
- CDPD-59138: [AUTOSYNC] Intermittent Delete root failed
- CDPD-59126: Seeing noexec permission on /tmp/liborg_apache_ratis_thirdparty_netty_transport_native_epoll_x86
- CDPD-58949: Import should not deduplicate schemas
- CDPD-58854: Few Ozone EC Distcp jobs are failing because pipeline limit has been reached
- CDPD-58848: Impala - Upgrade json-smart to 2.4.10 due to CVE-2023-1370
- CDPD-58652: Backport PHOENIX-6986 Add property to disable server merges for hinted uncovered indexes
- CDPD-58495: Ozone - Upgrade Netty Project to 4.1.94.Final due CVE-2023-34462
- CDPD-58220: ZDU | Getting java.lang.ClassNotFoundException: org.cloudera.log4j.redactor.RedactorAppender while starting ZEPPELIN
- CDPD-58029: [AUTOSYNC] Close open container immediately on ICR of unhealthy replica
- CDPD-58027: [AUTOSYNC] Fix Snapdiff output for key renames
- CDPD-58019: Ratis-Thirdparty - Bump guava to 32.0.0-jre
- CDPD-56724: Oozie web console is allowing access to list directories
- CDPD-56480: [AUTOSYNC] OmDBSnapshotInfoCodec.copyObject(..) does not follow the general contract of copy.
- CDPD-56456: Fix and backport PHOENIX-6961 Non-covered index failure with covered index fields
- CDPD-56176: Fix and backport PHOENIX-6910 Scans created during query compilation and execution against salted tables need to be more resilient
- CDPD-55637: [AUTOSYNC] ReplicationManager: Unhealthy replicas could block Ratis containers being recovered
- CDPD-55101: Invocation of Main class completed Message is skipped when LauncherSecurityManager calls system exit
- CDPD-55043: [AUTOSYNC] KeyDeleting service should not reclaim snapshot keys.
- CDPD-48979: Rotated Ranger KMS access logs aren't getting removed
- CDPD-30427: Fix custom ZooKeeper trust manager for FIPS
- [TSB 2023-703](#): Risk of Data Loss when using Hue S3 File Browser

Known issues in 7.1.9 CHF 1

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.1.9 CHF 1.

After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and 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, `ClassNotFoundException` for `org.cloudera.log4j.redactor.RedactorAppender` class was identified. The error message is harmless as the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

CDPD-68951: In 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 7.1.9 CHF3. However, the pre-existing (pre-upgrade) key's metadata cannot be changed.

None

When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(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 MPU. The problem is that 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.

CDPD-63665: The root cause is the netty native working directory configured by `org.apache.ratis.thirdparty.io.netty.native.workdir` with default value `${OZONE_HOME}/temp` which is `/opt/cloudera/parcels/CDH-7.1.9-1.cdh7.1.9.p1.47064069/lib/hadoop-ozone/` in the cluster is not available. The possible reason is that SCM instance does not have the permission to create directory under `hadoop-ozone`.

Add the configuration to java opts, with a new location that SCM process has write permission - `Dorg.apache.ratis.thirdparty.io.netty.native.workdir=${OZONE_HOME}/temp`.

OPSAPS-69539: CDP Runtime 7.1.9 does not support Oracle JDK 8u401 or OpenJDK 1.8.0_402 (8u402). Some services will fail to start. This can be a problem on RHEL 9.x as version 8u402 is the default OpenJDK 8 installed by the OS.

Workaround is to install an earlier version of JDK 8. For example Oracle `jdk-8u291` / `1.8.0_291`, or OpenJDK `8u292` / `1.8.0_292`.

CDPD-60839: KnoxShell clientside does not work with JDK17 due to incompatible groovy dependency.

The Knox service is working, however, KnoxShell is broken.

If you are using Knoxshell client, you must not upgrade to 7.1.9 CHF1.

CDPD-61524: Ozone Storage Container Manager fails to start on upgrading from CDP Private Cloud Base 7.1.6 to 7.1.9 CHF1. Also, if you have upgraded from CDP Private Cloud Base 7.1.6 to 7.1.7 or 7.1.8 and then to 7.1.9, the upgrade fails.

None. Cloudera recommends you to reach out to the Support before performing the upgrade to CDP Private Cloud Base 7.1.9.

CDPD-62254: Ozone is not supported on SLES15 with CHF1.

If your cluster has Ozone, Cloudera recommends you to not upgrade to 7.1.9 CHF1.

QAINFRA-18371: Conflict while installing libmysqlclient-devel on SLES 15

You may see an error such as the following while installing the mysql-devel and libmysqlclient-devel packages for setting up MariaDB as a backend database on SLES 15: File /usr/bin/mariadb_config from install of MariaDB-devel-<VERSION>.x86_64 conflicts with file from install of libmariadb-devel-3.1.21-150000.3.33.3.x86_64 (SLES Module Server Applications Updates)

While installing the mysql-devel and libmysqlclient-devel packages on SLES15, use the “--replac efiles” zypper switch or manually enter “yes” on the interactive pop-up that you see when the files are being overwritten.

CDPD-62464: Java process called by navatlas.sh tool fails on JDK-8 version

While running nav2atlas.sh script on OracleJDK 8 an error message is thrown and returns code 0 on an unsuccessful run.

You must install JDK-11 version on the host. Make sure not to put into the default path and JAVA_HOME. In a shell, set the JAVA_HOME to this location and run the nav2atlas.sh script.

CDPD-62834: Status of the deleted table is seen as ACTIVE in Atlas after the completion of navigator2atlas migration process

The status of the deleted table displays as ACTIVE.

None

CDPD-62837: During the navigator2atlas process, the hive_storagedesc is incomplete in Atlas

For the hive_storagedesc entity, some of the attributes are not getting populated.

None

CDPD-62935: If you are using the Knox Port Mapping feature, CDP Private Cloud Runtime 7.1.9 GA is not compatible with Cloudera Manager 7.11.3.2.

If you use Knox Port Mapping feature and want to upgrade Cloudera Manager to 7.11.3 CHF 1, then you must upgrade CDP Runtime to CDP 7.1.9 CHF 1.

COMPX-7493: YARN Tracking URL that is shown in the command line does not work when kinox is enabled

When Knox is configured for YARN, the Tracking URL printed in the command line of an YARN application such as spark-submit shows the direct URL instead of the Knox Gateway URL.

Upgrade CDP Runtime to CDP 7.1.9 CHF 2, and then you need to perform the following steps:

1. Open the Cloudera Manager Admin Console and go to the Knox service.
2. Click on the Knox Gateway Home URL.
3. Copy the YARN Resource Manager Web UI V2 URL from the Knox Gateway Home page.

For example, <https://knox-gateway.example.com:8443/gateway/cdp-proxy/yarnuiv2/>

4. Open the Cloudera Manager Admin Console and go to the YARN service.
5. Click the Configuration tab and search for resourcemanager_config_safety_valve.
6. Add the Resource Manager Advanced Configuration Snippet (Safety Valve) for yarn-site.xml property, and specify its value by using the YARN Resource Manager Web UI V2 URL, copied earlier, as follows:

```
Name: yarn.web-proxy.gateway.url
```

```
anager Web UI V2 URL> Value: <YARN Resource M
```

7. Enter a Reason for Change and then click Save Changes.
8. Restart YARN.

OPSAPS-69481: Some Kafka Connect metrics missing from Cloudera Manager due to conflicting definitions

The metric definitions for `kafka_connect_connector_task_metrics_batch_size_avg` and `kafka_connect_connector_task_metrics_batch_size_max` in recent Kafka CSDs conflict with previous definitions in other CSDs. This prevents Cloudera Manager from registering these metrics. It also results in SMM returning an error. The metrics also cannot be monitored in Cloudera Manager chart builder or queried using the Cloudera Manager API.

Contact Cloudera support for a workaround.

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet `atlas.authentication.method.file` is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

CDPD-79160: NPE while deleting BusinessMetadata

If business metadata is created without adding any applicable types, a `NullPointerException` is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

HBase shutdown can lead to inconsistencies in META

Cloudera Manager uses an incorrect shutdown command. This prevents graceful shutdown of the HBase service and forces Cloudera Manager to kill the processes instead. It can lead to inconsistencies in Meta.

Workaround: Run the following command instead of shutting down the HBase service using Cloudera Manager.

```
hbase master stop --shutDownCluster
```

The command output must end with Closing master protocol: `MasterService` phrase. You can verify the command execution by checking the master logs. The log must contain `Cluster shutdown requested of master=xxx` and the closing of regions. Upon successful execution, the `RegionServers` start shutting down.



Note: The command does not stop the *REST Server* and the *Thrift Server* role instances. You can safely shut down them from Cloudera Manager later.

If you find any inconsistencies, please contact Cloudera Support.

CDPD-76789: Creating tag with name description throws `java.lang.ClassCastException`

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails

Running `import-hive` command on RAZ enabled Cloudera Data Engineering cluster fails with error `"NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl"`.

CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in `getKafkaConsumer` ,`WakeupException: null`" is seen

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

UnsupportedClassVersionError

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.

CDPD-79099: Unable to import hive tables using import script in Public Cloud

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.

Technical Service Bulletins

TSB 2023-702: Potential wrong result for queries with date partition filter for clusters in GMT+ timezone

In Cloudera Data Platform (CDP) Private Cloud Base 7.1.7 Service Pack (SP) 2 Cumulative Hotfix (CHF) 11, a fix was introduced in Hive Metastore (HMS) to address a parsing issue with date strings. This fix caused a regression in Hive clusters where the HMS time zone is set ahead of GMT for the following combination of tables and queries: a table that is partitioned on a DATE column and a SELECT query on that table containing a WHERE clause filter on the same DATE column. For such queries, during the partition pruning phase, the date string would be converted to a date without timezone and compared with the partition value retrieved by HMS. This causes wrong results (0 rows) because the date values do not match.

The regression was identified in CDP Private Cloud Base 7.1.7 SP2 CHF14, but it exists in CHF11 through CHF16 as well as on certain versions of 7.1.8 and 7.1.9.

This issue does not affect clusters where the time zones are behind GMT. For example, if the time zone of the cluster is set USA/Los Angeles, which is 8 hours behind GMT, a date '2023-10-02' will remain as '2023-10-02' after converting to GMT (adding 8 hours). On the other hand, using Asia/Hong Kong time as an example, which is 8 hours ahead of GMT, the same date would become '2023-10-01' after converting to GMT (subtracting 8 hours), which leads to the wrong results.

Upstream JIRA

[HIVE-27760](#)

Knowledge article

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2023-702: Potential wrong result for queries with date partition filter for clusters in GMT+ timezone](#)

TSB 2024-749: Possible Ozone Snapshot Chain Corruption in 7.1.9.0

The Apache Ozone (Ozone) snapshots feature provides the ability to save consistent and immutable copies of the Ozone namespace. A chain of snapshots is tracked and maintained in the Ozone RocksDB that is utilized to calculate the difference between two snapshots required for the snapDiff API. A race condition on the snapshot delete path was found during testing. This condition might result in two consecutive snapshots pointing to the same parent, which breaks the snapshot chain.

Upstream JIRA

[HDDS-10524](#), [HDDS-10590](#), and [HDDS-9198](#)

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2024-749: Possible Ozone Snapshot Chain Corruption in 7.1.9.0](#)

On FIPS cluster, Knox-Impala connection is failing with SSL error code 5

The Knox-Impala connection is failing with SSL error code 5 on FIPS clusters. This prevents accessing Impala through Knox.

TSB 2025-820: Potential Data Integrity Issues Found in Ozone

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

Knowledge article

For the latest update on this issue see the corresponding Knowledge 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

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.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

Cumulative hotfix 2

You can review the list of cumulative hotfixes that were shipped for CDP Private Cloud Base version 7.1.9 CHF2.

Cloudera Runtime 7.1.9.3 (Cumulative Hotfix 2) download URL:

Parcel Repository Location

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.3/parcels/
```

Fixed issues in 7.1.9 CHF 2

Know more about the cumulative hotfixes 2 for 7.1.9. This cumulative hotfix was released on December 22, 2023.



Note: Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p3.48381316

- CDPD-63321: (HBASE-25643) The delayed FlushRegionEntry is replaced with the non-delayed one. The RegionServer periodically checks all the regions, if one is not flushed for a long time, then it creates a delayed FlushRegionEntry, the delay range is 0~300s. During the delay time, if many data are sent to the region, the flush can not be done immediately due to the existing one in regionsInQueue, and then the RegionTooBusyException occurs.
- CDPD-62205: A regression that was introduced in PHOENIX-6658 is fixed, which might affect users who still need to migrate to strongly consistent indexing. The regression issue affects the user's upgrade process and might also cause the old-styled indexes to go out of sync, which might require a full index rebuild.
- COMPX-15752: Queuemanager: configdiff validation error from 7.1.9 CHF1 -> CHF2
- COMPX-15602: Queuemanager loader spinner disappearing before UI is updated/re-rendered
- COMPX-15452: Verify that existing 7.1.9 customers migrated to Postgres can still use Postgres
- COMPX-15448: Testing normal functioning upgrades and fresh install clusters
- COMPX-15444: CPX changes to allow users to stay on H2 db - update database migration Info API
- COMPX-15443: Config Service Changes to disable QM migration using flag
- COMPX-15432: QueueManager flag to turn off database migration
- COMPX-15421: Fix broken UTs in CPX - ConfigServiceConnectorTest
- COMPX-15420: Update date format for DQS after switching to postgres DB
- COMPX-15375: CDPD - Upgrade Okhttp to 4.11.0 due to CVE-2023-0833 and CVE-2021-0341
- COMPX-15308: YARN-11578 Fix performance issue of permission check in verifyAndCreateRemoteLogDir
- COMPX-14855: Backporting YARN-11535 (Remove jackson-dataformat-yaml dependency)
- COMPX-14759: Queue Manager - Upgrade commons-configuration2 to 2.9.0 due to CVEs
- COMPX-14713: Backport YARN-11464 (TestFSQueueConverter#testAutoCreateV2FlagsInWeightMode has a missing dot before auto-queue-creation-v2.enabled for method call assertNoValueForQueues)
- COMPX-14064: Default ULF for dynamic queue template should be -1 in Weight mode
- COMPX-8329: Fix failing YARN Unit test:
TestYarnConfigurationFields.testCompareConfigurationClassAgainstXml - yarn.web-proxy.gateway.url
- COMPX-7493: YARN Tracking URL that is shown in the command line does not work when Knox is enabled
- COMPX-7247: Fix failing unit test:
org.apache.hadoop.yarn.server.nodemanager.containermanager.container.TestContainer.testKillOnNew
- CDPD-64790: Atlas build failure across release lines
- CDPD-64730: Oozie LauncherAM memory settings cannot be applied
- CDPD-64258: hive: Analyse compatibility report generated - between 7.1.9 CHF1 and CHF2
- CDPD-64240: CDPD-63145 causes regression in Orc
- CDPD-64133: The Oozie client should be able to handle Java 11+ related parameters
- CDPD-63799: Backport CDPD-45383 to 719 CHF and CDS 3.x CHFs
- CDPD-63780: Backport CDPD-42446 to 718 CHF and 719 CHF
- CDPD-63779: Oozie's spark actions are failing intermittently due to NPE
- CDPD-63769: [FIPS+JDK11] Solr health issue
- CDPD-63756: Backport CDPD-63231 to 7.1.8 CHF

- CDPD-63665: SCM Down:UnsatisfiedLinkError while enabling hdds.grpc.tls.enabled
- CDPD-63650: Fail early when encryption-at-rest init fails
- CDPD-63646: Fix file descriptor leak when encryption-at-rest is enabled
- CDPD-63602: Zeppelin - Upgrade jetty to 9.4.53/10.0.17/11.0.17 due to CVE-2023-40167, CVE-2023-36479, CVE-2023-41900, CVE-2023-36478 and CVE-2023-44487
- CDPD-63585: [AUTOSYNC] OM is getting stuck on snapshot creation if snapshot chain is corrupted
- CDPD-63520: Insecure direct object reference
- CDPD-63518: Both usersync/tagsync instances becoming active and syncing users/tags in the following scenario
- CDPD-63505: Backport IMPALA-12499 to 7.1.9 CHF
- CDPD-63504: Backport IMPALA-11068 to 7.1.9 CHF
- CDPD-63503: Backport IMPALA-12474 to 7.1.9 CHF
- CDPD-63502: Backport IMPALA-12492 to 7.1.9 CHF
- CDPD-63501: Backport IMPALA-12461 to 7.1.9 CHF
- CDPD-63500: Backport IMPALA-12460 to 7.1.9 CHF
- CDPD-63483: [FIPS+JDK11] - Quanta jobs failing with hadoop version cmd
- CDPD-63481: Backport IMPALA-12548 to CDH-7.1.9.x
- CDPD-63450: Backport HIVE-17350 to CDH-7.1.9.x
- CDPD-63438: java.lang.NullPointerException: at org.apache.hadoop.ozone.om.ratis.OzoneManagerRatisServer.getRaftLeaderId(OzoneManagerRatisServer.java:838)
- CDPD-63414: [AUTOSYNC] Legacy RM will not replicate all unhealthy containers when some are decommissioning
- CDPD-63393: [CDH-7.1.9 CHF2 CLONE] - AuthorizeOnlyWithChainedPolicies shows incorrect policy in Ranger audit when policy priority is equal
- CDPD-63351: [AUTOSYNC] Fix NPE in OMSnapshotPurgeRequest and exit loop early SnapshotDeletingService
- CDPD-63347: Backport CDPD-60975 to 7.1.8 CHFx , 7.1.9 CHFx
- CDPD-63321: Backport HBASE-25643 to 7.1.9 CHF2
- CDPD-63313: IMPALA-12542 test_query_cancel_created failed in ASAN build
- CDPD-63309: [UnitTest] testMaterializationLookup failure: timestamp mismatch
- CDPD-63308: Iceberg - Upgrade Netty Project to 4.1.100.Final due to CVE-2022-41881, CVE-2022-41915, CVE-2023-34462, CVE-2023-44487
- CDPD-63306: Zeppelin - Upgrade netty to 4.1.100.Final due to CVE-2023-44487
- CDPD-63302: Keytrustee-keyhsm - Upgrade Jetty to 9.4.53/10.0.17/11.0.17 due to CVE-2023-40167, CVE-2023-36479, CVE-2023-41900, CVE-2023-36478 and CVE-2023-44487
- CDPD-63301: SRM - Upgrade Jetty to 9.4.53/10.0.17/11.0.17 due to CVE-2023-40167, CVE-2023-36479, CVE-2023-41900, CVE-2023-36478 and CVE-2023-44487
- CDPD-63297: Knox - Upgrade Apache Santuario - xmlsec to 2.2.6/2.3.4/3.0.3 due to CVE-2023-44483
- CDPD-63294: Knox - Upgrade mysql-connector-j to 8.2.0 due to CVE-2023-22102
- CDPD-63290: Atlas - Upgrade amqp-client to 5.18.0+ due to CVE-2023-46120
- CDPD-63288: Schema Registry - Upgrade jose4j to 0.9.3 due to CVE-2023-31582
- CDPD-63283: IMPALA-12493 Impala Query cancelled while Analyzing or Compiling partially closes but query remains on Coordinator
- CDPD-63281: [AUTOSYNC] Backport HDDS-9550 to legacy RM (missing containers which are empty)
- CDPD-63244: [7.1.9 CHF2] - Not able to search using multiple user filter in access audit tab
- CDPD-63238: Parquet export fails with NoSuchMethodError
- CDPD-63202: Ranger kms is not getting started after cdp upgrade to 7.1.9
- CDPD-63180: Solr server unable to start after jetty upgrade to 9.4.53
- CDPD-63145: BytesColumnVector fails when the aggregate size is > 1gb
- CDPD-63139: [7.1.9 CHF2] [CLONE] - User name with comma split in old Ranger admin UI
- CDPD-63126: [AUTOSYNC] Snapdiff fails in case of key renames to deleted directories
- CDPD-63123: Sqoop build is taking 6 hours to complete

- CDPD-63120: [AUTOSYNC] DN import of container is not safe while replication
- CDPD-63119: [AUTOSYNC] Replication Manager could incorrectly use QUASI_CLOSED replicas as replication sources for CLOSED containers
- CDPD-63101: CLONE - Stored cross-site scripting on "Description" field under classification
- CDPD-63098: SMM - Upgrade Jetty to 9.4.53/10.0.17/11.0.17 due to CVE-2023-40167, CVE-2023-36479, CVE-2023-41900, CVE-2023-36478 and CVE-2023-44487
- CDPD-63061: Cruise Control - Upgrade org.json to 20231013+ due to CVE-2023-5072, CVE-2022-45688
- CDPD-63057: Cruise Control - Upgrade netty to 4.1.100.Final due to CVE-2023-44487, CVE-2023-34462
- CDPD-63032: [AUTOSYNC] Datanode should log Follower cannot close container at info level
- CDPD-63031: [AUTOSYNC] Log reason for not using a node at info level in SCMCommonPlacementPolicy
- CDPD-62992: [AUTOSYNC] Decommission should not wait on deleting containers
- CDPD-62982: [7.1.x] Ranger - Upgrade Json-Java to 20231013 due to CVE-2023-5072
- CDPD-62957: [AUTOSYNC] Container report shows missing containers when they actually appear empty
- CDPD-62935: [Analyze] [ST][Knox] test_knox_feature_topology_port_mapping tests fail
- CDPD-62927: Schemaregistry - Upgrade JSON-Java to 20231013 due to CVE-2023-5072
- CDPD-62926: Upgrade Json-Java to 20231013 due to CVE-2023-5072
- CDPD-62899: Backport HDDS-9432 to CDH-7.1.9.x
- CDPD-62868: Knox WaitForKnoxGatewayReadyToServeCommand reports http 500, service restart failure
- CDPD-62841: SRM - Upgrade Armeria to 1.26.0 due to CVE-2023-44487
- CDPD-62807: Backport HIVE-27558 to CDH-7.1.9.x
- CDPD-62800: ZooKeeper TLS/SSL support for Lucene-Solr
- CDPD-62791: NPE in SendContainerRequestHandler.deleteTarball
- CDPD-62788: Atlas [7.1.9 CHFx] - Upgrade netty to 4.1.100.Final due to CVE-2023-44487
- CDPD-62767: KnoxShell fails with Unsupported class file major version 61 error
- CDPD-62756: [AUTOSYNC] DataNode decommission retries for 300 times when invalid host or port is passed in the command
- CDPD-62741: org.apache.ratis.thirdparty.io.grpc.internal.ClientCallImpl\$ClientStreamListenerImpl \$1StreamClosed@7abc0029 java.lang.NullPointerException
- CDPD-62731: Backport HIVE-27772 to CDH-7.1.9.x
- CDPD-62724: HSTS header missing from unsecured API in Ranger Raz, Tagsync, Usersync
- CDPD-62721: HSTS header missing from unsecured API in Ranger Admin
- CDPD-62699: Harmonize jackson and jackson-databind for zookeeper
- CDPD-62666: Ignore used undeclared jetty dependency in phoenix-connectors
- CDPD-62657: Oozie applications might fail to start if default FIPS-related Java properties configured in Hadoop are not included in Oozie's Launcher AM's Java configuration.

In a FIPS environment, certain Java system properties must be passed to the Java Virtual Machine (JVM). In Hadoop, these properties are specified in configurations such as `yarn.app.mapreduce.am.admin-command-opts` or `yarn.app.mapreduce.am.command-opts`. Oozie must process these configurations to its Launcher AM application. Otherwise, the container will not start.

The process of injecting Oozie's Launcher AM configuration relies on several factors when assembling the launcher's Java system properties. For example, the default values of `oozie.launcher.override.javaopts` and `oozie.launcher.prepend.javaopts` direct Oozie to look up the values of `yarn.app.mapreduce.am.command-opts` and `yarn.app.mapreduce.am.admin-command-opts`, appending or prepending them to the Java system properties used in the Launcher AM container's start command. However, Oozie only searches for these values in the underlying action's configuration, the global Workflow configuration, `job.properties`, or the global action default configuration.

This fix extends the Launcher AM configuration injection by introducing the `oozie.LauncherConfigurationInjector.hadoop.search.properties` Oozie configuration. This property contains a comma-separated list of property names. The names listed in `oozie.LauncherConfigurationInjector.hadoop.search.properties` property allow Oozie to search not only Oozie-related configurations but also Hadoop configurations (`core-site.xml`, `hdfs-site.xml`, `mapred-site.xml`, `yarn-`

site.xml, or hadoop-site.xml). By default, `oozie.LauncherConfigurationInjector.hadoop.search.properties` is set to `yarn.app.mapreduce.am.admin-command-opts,yarn.app.mapreduce.am.command-opts`. So, if these properties are not set in any Workflow or Oozie configuration, but are present in a Hadoop configuration, Oozie uses the values from Hadoop. Thus, the default FIPS-related properties configured in Hadoop are automatically included in the Launcher AM's start command.

If you have already set a value for `yarn.app.mapreduce.am.admin-command-opts` or `yarn.app.mapreduce.am.command-opts` in `workflow.xml`, `job.properties`, or the global action default configuration, Oozie prioritizes those values and ignores the Hadoop configurations. Therefore, you must ensure that the FIPS-related system properties are included in the values you have configured.

- CDPD-62643: [AUTOSYNC] LegacyReplicationManager: Do not count unique origin nodes as over-replicated
- CDPD-62612: Backport ZOOKEEPER-4719 Use Bouncycastle jdk18on instead of jdk15on
- CDPD-62605: [7.1.9 CHF CLONE] - Upgrade Tomcat to 8.5.94+ (for CVE fixes) in all Ranger services
- CDPD-62591: Hue - Upgrade Tomcat to 9.0.81 due to CVE-2023-41080 and CVE-2023-44487
- CDPD-62586: Upgrade Tomcat to 8.5.94/9.0.81 due to CVE-2023-42794, CVE-2023-42795, CVE-2023-45648 and CVE-2023-44487
- CDPD-62567: Ratis-1886 AppendLog sleep fixed time cause significant drop in write throughput
- CDPD-62564: Atlas [7.1.9 CHFx] - Upgrade Okhttp to 4.11.0 due to CVE-2023-0833 and CVE-2021-0341
- CDPD-62557: Backport HIVE-27723 to CDH-7.1.9.x
- CDPD-62554: Backport HIVE-21100 to CDH-7.1.9.x
- CDPD-62538: [AUTOSYNC] A reformatted datanode node cannot be decommissioned
- CDPD-62513: SMM UI - Upgrade Node JS version to 20.8.1 due to multiple CVEs
- CDPD-62508: CDPD - Upgrade netty to 4.1.100.Final due to CVE-2023-44487 and CVE-2023-34462
- CDPD-62506: SMM - Upgrade netty to 4.1.100.Final due to CVE-2023-44487
- CDPD-62505: Kafka Connect Ext - Upgrade netty to 4.1.100.Final due to CVE-2023-44487
- CDPD-62504: Ratis thirdparty - Upgrade netty to 4.1.100.Final due to CVE-2023-44487
- CDPD-62503: Ozone - Upgrade netty to 4.1.100.Final due to CVE-2023-44487
- CDPD-62502: Ranger - Upgrade netty to 4.1.100.Final due to CVE-2023-44487
- CDPD-62501: Atlas - Upgrade netty to 4.1.100.Final due to CVE-2023-44487
- CDPD-62480: [AUTOSYNC] Non-blocking container statemachine cache
- CDPD-62456: Hive Acid Replication Support for Dell Powerscale - Backend Changes
- CDPD-62453: Backport HIVE-27760 to CDH-7.1.9.x
- CDPD-62448: Explicit handling of DIGEST-MD5 vs GSSAPI in quorum auth
- CDPD-62359: Backport PHOENIX-6994 Do not duplicate options specified in PHOENIX_QUERYSERVER_OPTS in queryserver.py
- CDPD-62348: Backport IMPALA-12462 to 7.1.9 CHF
- CDPD-62347: Backport IMPALA-8675 to 7.1.9 CHF
- CDPD-62312: Re-enable dependency harmonization for ZooKeeper
- CDPD-62297: Oozie unit tests do not clean up tens of GigaBytes of data causing UT container eviction
- CDPD-62264: Backport HIVE-27673 to CDH-7.1.9.x
- CDPD-62233: [snapshot] OM shuts down intermittently due to RocksDBException on createSnapshot request
- CDPD-62230: [snapshot] OM shutdown on RocksDB failure when performing distcp of snapshots
- CDPD-62224: Livy - Upgrade Okhttp to 4.11.0 due to CVE-2023-0833 and CVE-2021-0341
- CDPD-62222: Cruise Control - Upgrade Okhttp to 4.11.0 due to CVE-2023-0833 and CVE-2021-0341
- CDPD-62205: Backport PHOENIX-7057 Potential bug in MetadataEndpointImpl#updateIndexState.
- CDPD-62173: Merge HIVE-24530 on all CDP-PvC 7.1.[7-9] CHFx versions
- CDPD-62156: IMPALA-10860 Allow setting separate mem_limit for coordinators
- CDPD-62145: FIPS in Streaming with Java 11
- CDPD-62128: Using centralised version of snappy-java in Search
- CDPD-62126: Using centralised version of snappy-java in Solr
- CDPD-62125: Kafka - Upgrade snappy-java to 1.1.10.5 due to CVE-2023-43642
- CDPD-62063: Backport HIVE-27728 to CDP.

- CDPD-62059: AvroConnectTranslator should handle null values in fromConnectData method
- CDPD-62057: DefaultDispatch doesn't forward inbound request headers in case of requestType=OPTIONS
- CDPD-62046: Disable TestFanOutOneBlockAsyncDFSOutput
- CDPD-61986: Parcel impala-shell binaries won't work with non-standard Python 3 version
- CDPD-61951: Backport ZOOKEEPER-4674 TestReadOnlyClient.cc: Stop/start "normal" server in test setUp/tearDown
- CDPD-61917: Atlas - Upgrade Spring Security to 5.7.10/5.8.5/6.0.5/6.1.2 due to CVE-2023-34034 and CVE-2023-34035
- CDPD-61814: [7.1.9 CHF2] Implement best coding practices for validating user input
- CDPD-61798: Cannot drop unbounded range partitions in Kudu tables
- CDPD-61741: Backport HIVE-22613 to CDP.
- CDPD-61737: [AUTOSYNC] LegacyReplicationManager: Unhealthy replicas of a sufficiently replicated container can block decommissioning
- CDPD-61726: Backport Hive-27665 for CDH-7.1.9
- CDPD-61684: [AUTOSYNC] ReplicationManager: Ignore any Datanodes that are not in-service and healthy when finding unique origins
- CDPD-61625: Implement best coding practices for validating user input
- CDPD-61606: Potential dataloss from quick navigation during move op for S3 in Hue
- CDPD-61600: [AUTOSYNC] ReplicationManager: Handle all UNHEALTHY replicas of a CLOSING container
- CDPD-61589: Hue download from ABFS can return a corrupted file
- CDPD-61578: Impala - Upgrade Jetty to 9.4.53/10.0.17/11.0.17 due to CVE-2023-26048, CVE-2023-26049, CVE-2023-40167, CVE-2023-36479, CVE-2023-41900, CVE-2023-36478 and CVE-2023-44487
- CDPD-61577: CDPD - Upgrade Jetty to 9.4.53/10.0.17/11.0.17 due to CVE-2023-40167, CVE-2023-36479, CVE-2023-41900, CVE-2023-26048, CVE-2023-26049, CVE-2023-36478 and CVE-2023-44487
- CDPD-61535: [AUTOSYNC] Handle all UNHEALTHY replicas of a CLOSING container
- CDPD-61507: Atlas [7.1.9 CHFx] - Upgrade Okio to 3.4.0 due to CVE-2023-3635
- CDPD-61503: Atlas [7.1.9 CHFx] - Upgrade Apache Ivy to 2.5.2 due to CVE-2022-46751
- CDPD-61495: [AUTOSYNC] compactionLogTable to store compaction information
- CDPD-61489: [AUTOSYNC] LegacyReplicationManager: Unhealthy replicas of a sufficiently replicated container can block decommissioning
- CDPD-61385: [AUTOSYNC] LOG improvement when downloading container fails from DN
- CDPD-61380: [AUTOSYNC] Avoid creating Managed objects per request to avoid the finalizer cost
- CDPD-61248: [7.1.9 CLONE] - RangerKafkaAuditHandler broken and multiple authorizations audited in CDP 7.1.8
- CDPD-61172: Find out how to fix hbase-indexer cli for Zookeeper SSL
- CDPD-61132: API compatibility Whitelist for hive
- CDPD-61084: [7.1.x] - [FIPS + JDK11] Ranger ChangePasswordUtil fails when CM comes up with the addition of bctls.jar
- CDPD-61061: [AUTOSYNC] Ozone cli command to get container info should deal with empty values for --json
- CDPD-60946: IMPALA-12413 Make Iceberg tables created by Trino compatible with Impala
- CDPD-60846: ZooKeeper TLS/SSL support for Hive-Solr
- CDPD-60646: [AUTOSYNC] Snapshot chain corruption should not fail OM restart
- CDPD-60338: HIVE-27669: [HiveAcidReplication] Hive Acid CTAS fails incremental if no of rows inserted is > INT_MAX
- CDPD-60006: Backport HIVE-22489, HIVE-24883 and HIVE-25410 issues to fix java.lang.ClassCastException in join on array column
- CDPD-59847: Zeppelin - Upgrade jackrabbit-webdav to 2.21.18 due to CVE-2023-37895
- CDPD-59846: Upgrade jackrabbit-webdav to 2.21.18 due to CVE-2023-37895
- CDPD-59842: SRM - Upgrade Armeria to 1.24.3 due to CVE-2023-38493
- CDPD-59673: During discovery if cm is not reachable and throws SocketException then retry is not happening
- CDPD-59579: Upgrade Spring Security to 5.7.10/5.8.5/6.0.5/6.1.2 due to CVE-2023-34034 and CVE-2023-34035

- CDPD-59481: [UnitTest] testConnectionRetryExceptionListener fails w/ BindException: Address already in use
- CDPD-59365: CDPD - Upgrade Shiro to 1.12.0 due to CVE-2023-34478
- CDPD-58823: jwks.json doesn't have double quotes which makes json invalid
- CDPD-58290: [AUTOSYNC] SST files are missing on optimized snapDiff path.
- CDPD-58171: IMPALA-12245 TestWebPage::test_query_progress is flaky
- CDPD-57125: HIVE-21213: Acid table bootstrap replication needs to handle directory created by compaction with txn id
- CDPD-56486: [Spark] Ozone delete key failed error during Spark job completion
- CDPD-52462: Race condition in getFileStatus causes flaky testObjectStoreCreateWithO3fs
- CDPD-52433: [Snapshot] Use RocksDB to persist compaction log
- CDPD-51430: Create Container failed using a disk which is full
- CDPD-50915: Oozie shouldn't ignore hive-site.xml on host if no hive-site is on Spark share lib
- CDPD-50443: Upgrade CM API usage for discovery
- CDPD-44719: Spark Atlas Connector - Update log4j to reload4j
- CDPD-42384: Spark Atlas Connector - Upgrade Data Mapper for Jackson to 1.9.16-TALEND due to high CVEs
- CDPD-41138: Impala - Upgrade jdom to 2.0.6.1 due to CVE-2021-33813
- CDPD-35383: Add entry in replication_metrics table for skipped/failed replication.
- CDPD-18153: ZooKeeper TLS/SSL support for solr-upgrade.sh
- CDPD-8443: RemoteException when moving a file from scratchdir to a directory in encryption zone
- [TSB 2023-702](#): Potential wrong result for queries with date partition filter for clusters in GMT+ timezone
- [TSB 2023-704](#): File corruption when downloading files larger than 1 MB from ABFS with Hue File Browser

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

- CVE-2010-5312
- CVE-2011-4969
- CVE-2012-6708
- CVE-2015-0897
- CVE-2016-7103
- CVE-2020-15522
- CVE-2020-26870
- CVE-2020-28458
- CVE-2020-7656
- CVE-2021-23445
- CVE-2021-33813
- CVE-2021-41182
- CVE-2021-41183
- CVE-2021-41184
- CVE-2022-2047
- CVE-2022-2048
- CVE-2022-31160
- CVE-2023-34034
- CVE-2023-34035
- CVE-2023-34478
- CVE-2023-36052
- CVE-2023-37895
- CVE-2023-38493
- CVE-2023-42794
- CVE-2023-42795
- CVE-2023-45648
- CVE-2023-4586

- CVE-2023-46120
- CVE-2023-5072

Known issues in 7.1.9 CHF 2

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.1.9 CHF 2.

After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and 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, `ClassNotFoundException` for `org.cloudera.log4j.redactor.RedactorAppender` class was identified. The error message is harmless as the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

CDPD-68951: In 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 7.1.9 CHF3. However, the pre-existing (pre-upgrade) key's metadata cannot be changed.

None

When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(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 MPU. The problem is that 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.

OPSAPS-69539: CDP Runtime 7.1.9 does not support Oracle JDK 8u401 or OpenJDK 1.8.0_402 (8u402). Some services will fail to start. This can be a problem on RHEL 9.x as version 8u402 is the default OpenJDK 8 installed by the OS.

Workaround is to install an earlier version of JDK 8. For example Oracle jdk-8u291 / 1.8.0_291, or OpenJDK 8u292 / 1.8.0_292.

Unsupported or uncertified components integration with Ozone:

Currently, Livy and Zeppelin are not certified for ozone integration.

When you upgrade from Cloudera Runtime 7.1.7 to 7.1.9, Livy cannot access Ozone FS because of renaming the ozone-file-system-hadoop3-*.jar file in Cloudera Runtime 7.1.7 without a corresponding update in Livy's configuration.

You must manually add ozone-file-system-hadoop3-*.jar to the Livy classpath.



Note: This issue is resolved for the Spark and Zeppelin components in Cloudera Runtime 7.1.8.

A fresh install of 7.1.9 or its CHF's does not allow user to bypass the Setup Database screen for YARN Queue Manager

YARN Queue Manager in Cloudera Data Platform (CDP) Private Cloud Base 7.1.9 CHF 2 does not require you to install a PostGres database, therefore users should not see the Setup Database screen and should be able to skip the Setup Database screen. With this known issue, users who are conducting a fresh install of 7.1.9 or its CHF's are not able to bypass the Setup Database screen as expected.

1. When conducting a fresh install of YARN Queue Manager in 7.1.9 or its CHF's, you must ensure that you have both CDP and Cloudera Manager upgraded to 7.1.9 CHF 2.
2. When you reach the Setup Database screen in the Cloudera Manager installation wizard for Queue Manager, enter any dummy values for the following fields:
 - a. Database name: configstore
 - b. Database Username: dbuser
 - c. Database Password: dbpassword
 - d. Database Hostname: localhost

YARN Queue Manager will not connect to PostGres with the above details and will fall back to the embedded database.

3. Run the following script command in a browser console to enable the Continue button:

```
document.querySelector('.btn.next').removeAttribute('disabled');
```

4. Click Continue and proceed with the YARN Queue Manager installation.
5. After installation is complete, SSH into the host that has Queue Manager installed, and run this command: `sed -i 's/migrationCompleted=true/migrationCompleted=false/' /var/lib/hadoop-yarn/migration.properties`



Note: Enable Queue Manager in the YARN configurations, and restart YARN.

6. Restart YARN Queue Manager.

CDPD-61524: Ozone Storage Container Manager fails to start on upgrading from CDP Private Cloud Base 7.1.6 to 7.1.9 CHF1. Also, if you have upgraded from CDP Private Cloud Base 7.1.6 to 7.1.7 or 7.1.8 and then to 7.1.9, the upgrade fails.

None. Cloudera recommends you to reach out to the Support before performing the upgrade to CDP Private Cloud Base 7.1.9.

CDPD-62254: Ozone is not supported on SLES15 with CHF1.

If your cluster has Ozone, Cloudera recommends you to not upgrade to 7.1.9 CHF1.

QAINFRA-18371: Conflict while installing libmysqlclient-devel on SLES 15

You may see an error such as the following while installing the mysql-devel and libmysqlclient-devel packages for setting up MariaDB as a backend database on SLES 15: File /usr/bin/

mariadb_config from install of MariaDB-devel-<VERSION>.x86_64 conflicts with file from install of libmariadb-devel-3.1.21-150000.3.33.3.x86_64 (SLES Module Server Applications Updates)

While installing the mysql-devel and libmysqlclient-devel packages on SLES15, use the “--replac efiles” zypper switch or manually enter “yes” on the interactive pop-up that you see when the files are being overwritten.

CDPD-62464: Java process called by navatlas.sh tool fails on JDK-8 version

While running nav2atlas.sh script on OracleJDK 8 an error message is thrown and returns code 0 on an unsuccessful run.

You must install JDK-11 version on the host. Make sure not to put into the default path and JAVA_HOME. In a shell, set the JAVA_HOME to this location and run the nav2atlas.sh script.

CDPD-62834: Status of the deleted table is seen as ACTIVE in Atlas after the completion of navigator2atlas migration process

The status of the deleted table displays as ACTIVE.

None

CDPD-62837: During the navigator2atlas process, the hive_storagedesc is incomplete in Atlas

For the hive_storagedesc entity, some of the attributes are not getting populated.

None

CDPD-63690: RuntimeException encountered when generating snapshotDiff report between 2 snapshots

When snapshot feature is enabled, KeyDeletingService, SSTFilteringService and SnapDiff thread fall into a deadlock when accessing Snapshot Cache.

Restart the Ozone Manager.

CDPD-63874: Changing OM service ID can cause OM startup failure and reverting to the original service ID also causes OM to go to a bad state.

OM service ID should not be changed as it is used to construct the raft group directory. If the change is made for a fresh install, delete the original data and metadata directories.

CDPD-64238: Snapshot diff request failing when setting ozone.om.snapshot.db.max.open.files=-1

When snapshot feature is enabled, KeyDeletingService, SSTFilteringService and SnapDiff thread fall into a deadlock when accessing Snapshot Cache.

Restart the Ozone Manager.

OPSAPS-69481: Some Kafka Connect metrics missing from Cloudera Manager due to conflicting definitions

The metric definitions for kafka_connect_connector_task_metrics_batch_size_avg and kafka_connect_connector_task_metrics_batch_size_max in recent Kafka CSDs conflict with previous definitions in other CSDs. This prevents Cloudera Manager from registering these metrics. It also results in SMM returning an error. The metrics also cannot be monitored in Cloudera Manager chart builder or queried using the Cloudera Manager API.

Contact Cloudera support for a workaround.

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

CDPD-79160: NPE while deleting BusinessMetadata

If business metadata is created without adding any applicable types, a NullPointerException is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

CDPD-76789: Creating tag with name description throws java.lang.ClassCastException

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

UnsupportedClassVersionError

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.

CDPD-79099: Unable to import hive tables using import script in Public Cloud

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.

Technical Service Bulletins**TSB 2024-749: Possible Ozone Snapshot Chain Corruption in 7.1.9.0**

The Apache Ozone (Ozone) snapshots feature provides the ability to save consistent and immutable copies of the Ozone namespace. A chain of snapshots is tracked and maintained in the Ozone RocksDB that is utilized to calculate the difference between two snapshots required for the snapDiff API. A race condition on the snapshot delete path was found during testing. This condition might result in two consecutive snapshots pointing to the same parent, which breaks the snapshot chain.

Upstream JIRA

[HDDS-10524](#), [HDDS-10590](#), and [HDDS-9198](#)

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2024-749: Possible Ozone Snapshot Chain Corruption in 7.1.9.0](#)

TSB 2025-820: Potential Data Integrity Issues Found in Ozone

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

Knowledge article

For the latest update on this issue see the corresponding Knowledge 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

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.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

YARN Queue Manager

Learn about the new features for YARN Queue Manager in Cloudera Runtime 7.1.9 cumulative hotfix (CHF) 2.

Database requirement changes

This release allows you to continue using your embedded database and no longer requires you to use a PostgreSQL database as required for Cloudera Data Platform (CDP) 7.1.9 and CDP 7.1.9 cumulative hotfix (CHF) 1. The option to migrate from the embedded database to PostgreSQL will be coming in a later Service Pack.

Java heap size can now be configured

You can now customize Java heap size in YARN Queue Manager. Although the default for this setting should be valid in most deployment scenarios, you have the option to update the setting only if a given cluster has run into memory-management issues, otherwise, the settings can remain.

Cumulative hotfix 3

You can review the list of cumulative hotfixes that were shipped for CDP Private Cloud Base version 7.1.9 CHF3.

Cloudera Runtime 7.1.9.4 (Cumulative Hotfix 3) download URL:**Parcel Repository Location**

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.4/parcels/
```

JDK supported information

OpenJDK 8 / Oracle JDK 8 in any CDP 7.1.9 and Cloudera Manager 7.11.3 version, or OpenJDK 11 / Oracle JDK 11 starting in CDP 7.1.9 CHF 3 and Cloudera Manager 7.11.3 CHF 3 versions and higher, as designated with p3 or .3 suffixes.

Fixed issues in 7.1.9 CHF 3

Know more about the cumulative hotfixes 3 for 7.1.9. This cumulative hotfix was released on February 23, 2024.



Note: Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p4.50495721

- [TSB 2024-724](#) - Issue with Apache Ozone non-HA SCMs during installation/upgrade when security tokens are enabled.
- CDPD-59477: CreateFile is not setting isFile flag in OmKeyInfo. This issue is resolved.
- KT-7508: Keytrustee-keyhsm - Upgrade Bouncy Castle to 1.74 due to CVE-2023-33202 and CVE-2023-33201
- KT-7506: [FIPS+JDK11] KeyTrustee Server fails with openssl command error on outputting keys and certificates
- COMPX-15869: [7.1.7 SP3, 7.1.9 CHF3, 7.2.19] - Queue Manager: Upgrade Okio to 3.4.0 due to CVE-2023-3635
- COMPX-15833: Use centralized jackson version in QueueManager
- COMPX-15798: Backport YARN-11630 (Passing admin Java options to container localizers)
- COMPX-15737: QM - Upgrade Bouncy Castle to 1.74 due to CVE-2023-33202 and CVE-2023-33201
- COMPX-15347: QM - Upgrade Plexus-utils to 3.3.1+ due to CVE-2022-4244 and CVE-2022-4245
- COMPX-15205: QM - Upgrade wiremock-jre8 to 2.35.1 due to CVE-2023-41327 and CVE-2023-41329
- COMPX-15161: Conversion from absolute to relative mode fails for low memory values
- COMPX-14794: CPX - Upgrade moment.js to 2.29.4 due to CVE-2022-24785, CVE-2022-31129
- COMPX-11123: Queue Manager - Upgrade Commons IO to 2.11.0/20030203.000550 due to medium CVEs
- COMPX-7242: Fix failing unit tests:
org.apache.hadoop.yarn.server.resourcemanager.scheduler.fair.TestContinuousScheduling
- COMPX-7241: Fix failing unit test:
org.apache.hadoop.yarn.client.api.impl.TestAMRMProxy.testAMRMProxyTokenRenewal
- COMPX-6271: Fix failing unit test:
org.apache.hadoop.yarn.server.resourcemanager.webapp.TestRMWebServicesNodesScaling.testClusterScalingInfoJson
- COMPX-6254: Fix failing unit tests: org.apache.hadoop.yarn.client.api.impl.TestNMClient
- CDPD-66138: [7.1.9] Fix ranger kafka plugin junit test failures
- CDPD-65548: Multi-master config fails with check failed
- CDPD-65493: #1396168-P1-hue-7.3.0.0-107 Build Error
- CDPD-65475: Backport "Too many "Failed to accept allocation proposal" because of wrong Headroom check for DRF" to CDH-7.1.9.x
- CDPD-65397: [AUTOSYNC] UNHEALTHY replicas of QUASI_CLOSED container with unique origins should be handled during decommission
- CDPD-65341: [AUTOSYNC] Snapshot read calls are failing due to SnapshotCache's inconsistency
- CDPD-65316: Backport HIVE-27919 to CDH-7.1.9.x
- CDPD-65315: Backport HIVE-27658 to CDH-7.1.9.x
- CDPD-65292: [7.1.9.CHFx]Atlas UI: Change the alignment of the Download Search button on the Classic UI search page.
- CDPD-65284: Parcel impala-shell won't work with Python 3.8 on SLES 15

- CDPD-65279: Phoenix - Upgrade Bouncy Castle to 1.70 due to medium CVEs
- CDPD-65277: Backport IMPALA-12595 to 7.1.9CHF3
- CDPD-65256: Backport TEZ-3972 to CDH-7.1.9.x
- CDPD-65243: Backport IMPALA-12683 to 7.1.9 CHF
- CDPD-65206: Backport IMPALA-12577 to 7.1.9 CHF
- CDPD-65169: [AUTOSYNC] Add number of datanodes, total capacity/used space to SCMNodeMetrics
- CDPD-65159: [AUTOSYNC] Snapshot: 'ozone fs -ls' on '.snapshot' dir of a bucket should list only active snapshots
- CDPD-65049: HTTP security headers are missing from Oozie response
- CDPD-65048: Backport HIVE-26208 to CDH-7.1.9.x
- CDPD-65043: Livy - [7.1.9 CHFx] Upgrade datatables to 1.10.23+ due to CVE-2020-28458
- CDPD-65039: Backport HDDS-8822. [S3G] Improve list performance in LEGACY bucket
- CDPD-65038: Backport HDDS-8011. IllegalArgumentException logged for invalid user-defined metadata
- CDPD-65036: Backport HDDS-9627. Reset RaftPeer priorities after transfer leadership
- CDPD-65035: Backport HDDS-9314. create-bucket on an existing bucket for s3g does not fail
- CDPD-65032: Backport HDDS-9708. Fix unit tests to reuse DispatcherContext
- CDPD-65031: Backport HDDS-9697 ContainerStateMachine.applyTransaction(..) should not validate token again
- CDPD-65013: CDPD - Upgrade Apache Shiro to 1.13.0 due to CVE-2023-46750
- CDPD-65012: Upgrade Apache Shiro to 1.13.0 due to CVE-2023-46750
- CDPD-65003: Centralize missing dependencies of Zeppelin to CDPD
- CDPD-64948: Temporary fix compile error from HDDS-9709
- CDPD-64919: Backport HIVE-24858 to CDH-7.1.9.x
- CDPD-64905: Backport IMPALA-12589 to active branches
- CDPD-64800: Classic UI - Security zone form not populate resources value properly while creating and editing zone form.
- CDPD-64798: [7.2.18.0 & 7.1.9 CHF3] - Keep the LDAP usersync details popup names same as the backbone js names
- CDPD-64747: Use centralized gson version in Zeppelin
- CDPD-64736: [7.2.18.0 & 7.1.9 CHF3] Fix to use correct service for resource lookup API in security zone
- CDPD-64734: Use centralized nimbus-jose-jwt version in Cruise Control
- CDPD-64726: 71x backport - Slowness / broadcast timeout issues due to SPARK-33290: REFRESH TABLE should invalidate cache even though the table itself may not be cached (Spark 2.4.8)
- CDPD-64720: Replace PHOENIX-6721 with the upstream version.
- CDPD-64707: hue build failure in centos7
- CDPD-64665: [MANUAL SYNC] Refine certificate renewer service to avoid it scheduled ahead of time
- CDPD-64648: Backport the versionless bigtop-new gerrits into 7.1.8 and 7.1.9
- CDPD-64627: [7.1.x]- Ranger - Upgrade Apache Derby to 10.17.1.0 due to CVE-2022-46337
- CDPD-64584: [7.1.9 CHF3] Upgrade Tomcat to 8.5.96 (for CVE fixes) in all Ranger services
- CDPD-64580: Allow ozone admin container info to list multiple containers
- CDPD-64566: Backport PHOENIX-7148 Use getColumnLabel Instead of getColumnName in QueryServerBasicsIT
- CDPD-64562: [AUTOSYNC] Decommission: Admin monitor should call RM.checkContainerState to check for under-replication
- CDPD-64550: Backport PHOENIX-7143 Detect JVM version and add the necessary flags in PQS startup script (phoenix query server repo)
- CDPD-64539: Postpone CM configuration change monitoring until the Knox GW is up&running
- CDPD-64527: Unable to place replicas using range aware logic with multiple locations
- CDPD-64517: Kafka connect S3 connector failing with AWS error
- CDPD-64480: Set name field with qualifiedName for impala_process and impala_process_execution
- CDPD-64478: Optimize Relationship Edge fetch

- CDPD-64450: Backport PHOENIX-7143 Detect JVM version and add the necessary flags in PQS startup script (phoenix repo)
- CDPD-64449: Backport the JVM module options from branch-2.4 HEAD
- CDPD-64444: HWC Full GC : Stack Overflow Error fails cdh-7.1.9.x builds
- CDPD-64427: LDAP group import/sync fails for "memberUid"
- CDPD-64425: [FIPS+JDK11] Intermittent Kafka connection issues during installation
- CDPD-64419: OM nodes went down due to OOM, possible memory leak
- CDPD-64398: [AUTOSYNC] OM/DN startup failure with non-HA SCM for secret manager not initialized
- CDPD-64376: Oozie's Spark and Spark3 option parser does not respect Java arguments starting with '--'
- CDPD-64372: OzoneManager - isDBUpdateSuccess flag not being set at OM client causes incorrect behaviour at Recon and failed to Recover in case of rocksDB exception
- CDPD-64364: OM/DN startup failure with non-HA SCM for secret manager not initialized
- CDPD-64347: Extend Java opts for Livy to support JDK17 + Isilon
- CDPD-64335: Zeppelin - Upgrade Bouncy Castle to 1.74 due to CVE-2023-33202 and CVE-2023-33201
- CDPD-64302: Remove Derby dependency in Solr.
- CDPD-64281: Backport HIVE-26802: Create qtest running QB compaction queries for ACID, insert-only and clustered tables
- CDPD-64272: Atlas [7.1.9 CHFx] - Upgrade reactor-netty to 1.0.39/1.1.13 due to CVE-2023-34062
- CDPD-64243: Backport HIVE-27643: Exclude compaction queries from ranger policies
- CDPD-64229: Impala - Upgrade Apache Derby to 10.17.1.0 due to CVE-2022-46337
- CDPD-64225: Sqoop - Upgrade Apache Derby to 10.17.1.0 due to CVE-2022-46337
- CDPD-64192: [7.1.9] - Atlas Server side Ignore and Prune patterns doesn't work
- CDPD-64184: Ozone resource lookup is not working due to "Service ID specified does not match with ozone.om.service.ids defined in the configuration."
- CDPD-64159: [7.1.9.x] - Ranger policy delta issue causing intermittent permission deny for Hive and HDFS services
- CDPD-64129: Backport HIVE-25684 to CDH-7.1.9.x
- CDPD-64123: Schema Registry - Upgrade Netty Project to 4.1.100.Final due to CVE-2023-44487
- CDPD-64122: CDPD - Upgrade aws-java-sdk-bundle to 1.12.599 due to CVE-2023-44487
- CDPD-64115: Impala build failure for 7.1.9.1
- CDPD-64114: Atlas - Upgrade reactor-netty to 1.0.39/1.1.13 due to CVE-2023-34062 and CVE-2023-34054
- CDPD-64040: Remove the CDP versions from Spark 2 deprecation message
- CDPD-64037: [7.1.9 CHF3] - "Select All permissions for all components." checkbox missing in tag based policy permission popup
- CDPD-64032: [7.1.9.CHFx] Atlas UI Basic Searching result sorting option not available on all Columns.
- CDPD-64019: [AUTOSYNC] Provide a flag to skip the native_rocksdb_tool loading
- CDPD-64007: Backport HIVE-27885 on CDP branches
- CDPD-64000: [AUTOSYNC] Datanode Write performance degradation
- CDPD-63962: [AUTOSYNC] Over Replication Check of all UNHEALTHY replicas is broken
- CDPD-63956: [AUTOSYNC] SCM's FinalizationStateManager#finalizeLayoutFeature Ratis call should be idempotent
- CDPD-63947: [AUTOSYNC] Disable rocksDB cache for snapshot
- CDPD-63915: Sqoop Teradata export fails if source table is empty
- CDPD-63874: Changing the Ozone service Id makes the cluster[OM] state irrecoverable
- CDPD-63849: [AUTOSYNC] Legacy Replication Manager should consider that UNHEALTHY replicas might be decommissioning
- CDPD-63841: [AUTOSYNC] OM fails with Snapshot chain corruption during SnapshotPurge
- CDPD-63839: [AUTOSYNC] Incorrect sorting order in RatisOverReplicationHandler
- CDPD-63837: [AUTOSYNC] Infinite loop in ReconUtils.nextClosestPowerIndexofTwo()
- CDPD-63835: Backport HIVE-27679 on all CDP-PvC 7.1.[7-9] CHFx versions
- CDPD-63804: [AUTOSYNC] SCM WebUI incorrectly renders DN links

- CDPD-63783: [AUTOSYNC] Provide API to check a container via Replication Manager
- CDPD-63734: [AUTOSYNC] NO_REPLICA_FOUND should trigger a OM pipeline cache refresh
- CDPD-63733: [AUTOSYNC] Missing snapshot entries list Snapshot under a bucket API
- CDPD-63724: Add spark-sql-kafka to Oozie Spark/Spark3 share libs
- CDPD-63723: Sqoop should determine files as Parquet by PAR1 in header
- CDPD-63692: In Rms- s3, db level access write permission mapping config is not working
- CDPD-63623: [UnitTest] Some Oozie units are failing due to HCat related NPE
- CDPD-63606: Datanodes do not Retry Pipeline Close Commands for SCM
- CDPD-63600: HDFS Authorizer changes to take advantage of support for multiple access-types in the Ranger Access Request (RANGER-4007)
- CDPD-63588: Do not show empty containers as missing in Recon UI
- CDPD-63574: disableLoadBalancingForUserAgents cannot be set
- CDPD-63571: TestIcebergTable.test_hive_external_forbidden fails on 7.1.9 builds
- CDPD-63570: TestIcebergTable.test_iceberg_negative fails on 7.1.9 builds
- CDPD-63553: [AUTOSYNC] Containers belonging to out of service nodes, are counted as mis-replicated
- CDPD-63527: [AUTOSYNC] Read from non-datanode host does not consider topology
- CDPD-63523: [AUTOSYNC] Topology level is not set in datanode object for distance calculation
- CDPD-63440: CLONE - UI: Enum type Business metadata attribute shows incorrect data when specific string is in attribute name.
- CDPD-63371: [AUTOSYNC] Parallel loading datanode volume db store
- CDPD-63326: Fix CVE-2023-36877 Apache Oozie Spoofing Vulnerability
- CDPD-63291: Search - Upgrade amqp-client to 5.18.0+ due to CVE-2023-46120
- CDPD-63287: Solr - Upgrade jose4j to 0.9.3 due to CVE-2023-31582
- CDPD-63286: Upgrade jose4j to 0.9.3 due to CVE-2023-31582
- CDPD-63276: [AUTOSYNC] Overwrite file by multipart upload, saving wrong ReplicationConfig in KeyInfo
- CDPD-63124: Newly added Kudu master couldn't start on custom kerberos cluster
- CDPD-63118: [AUTOSYNC] Replication Manager: Save UNHEALTHY replicas with highest BCSID for a QUASI_CLOSED container
- CDPD-63117: [AUTOSYNC] Replication Manager: Do not count unique origin nodes as over-replicated
- CDPD-63116: [AUTOSYNC] Make the number of containers logged configurable in DatanodeAdminMonitorImpl
- CDPD-63030: [AUTOSYNC] Snapshot diff job failed due to Metrics source OmSnapshotMetrics already exists
- CDPD-62943: [AUTOSYNC] NPE in OMDBCCheckpointServlet with ozone.om.ratis.enable=false
- CDPD-62942: [AUTOSYNC] Fix possible deadlock during shutdown in OzoneDelegationTokenSecretManager
- CDPD-62881: [AUTOSYNC] Recon - NPE in handling deleteKey event in NSSummaryFSO task
- CDPD-62826: [AUTOSYNC] Two S3G instances writing the same key may cause data loss in case of an exception.
- CDPD-62719: Datanode should not need to download existing container
- CDPD-62540: [AUTOSYNC] Pipeline close doesn't wait for containers to be closed
- CDPD-62464: Java process called by nav2atlas.sh tool fails on JDK8
- CDPD-62436: [AUTOSYNC] S3 default GRPC transport doesn't utilize enough parallelism on OM server-side
- CDPD-62429: [AUTOSYNC] TypedTable prefix iterator may leak CodecBuffer
- CDPD-62276: [AUTOSYNC] 'java.lang.UnsatisfiedLinkError' when trying to read RocksDB with 'ozone debug ldb'
- CDPD-62095: Backport HIVE-27525 to CDP
- CDPD-61962: [AUTOSYNC] Reduce the number of system calls when DN writes a key
- CDPD-61913: [AUTOSYNC] Avoid copying ByteString in ByteStringCodec
- CDPD-61754: Unknown container from datanode in Recon
- CDPD-61742: Test failure: org.apache.spark.sql.hive.execution.HiveTableScanSuite.Spark-4077: timestamp query for null value

- CDPD-61692: [AUTOSYNC] ReplicationManager: Ignore any Datanodes that are not in-service and healthy when finding unique origins
- CDPD-61659: [7.1.9 CHF3] Options for permissions pop up for Knox policies are not the same in Backbone UI and React JS
- CDPD-61626: [7.1.9 CHF3] - Keep the usersync details popup names same as the backbone js names
- CDPD-61539: [AUTOSYNC] Better datanode exclude list handling for long-lived clients
- CDPD-61492: [AUTOSYNC] Fix comparison logic for SCMContainerPlacementCapacity.
- CDPD-61425: [AUTOSYNC] Speed up TestStorageContainerManagerHA
- CDPD-61336: Canary build failing with upstream Ozone master branch
- CDPD-61251: Zookeeper - Upgrade jackson-databind to 2.13.4.1+ due to CVE-2022-42003, CVE-2022-42004
- CDPD-61068: [AUTOSYNC] Write performance degradation
- CDPD-60989: Upgrade Ozone upstream version for 7.1.9 release
- CDPD-60892: [AUTOSYNC] [FSO] S3A compatibility - dfs -put creates dir and a file
- CDPD-60882: [AUTOSYNC] Poor S3G read performance
- CDPD-60664: [AUTOSYNC] Snapshot Bootstrap creates incorrect hard links.
- CDPD-60592: [AUTOSYNC] OzoneManager: NPE on ACLs check in case of multipart upload to EC-bucket
- CDPD-60367: [AUTOSYNC] Invalidate snapshot cache once snapshot gets purged
- CDPD-60242: [AUTOSYNC] [Hsync] moves blocks to deleted table on final commit
- CDPD-60126: [AUTOSYNC] Potential data loss with HSync due to deletedTable entry having the same block as keyTable entry's
- CDPD-60070: [AUTOSYNC] Use sequence ID for certificate serial ID
- CDPD-59781: [AUTOSYNC] Bucket replication type is ignored when uploading files via S3G
- CDPD-59477: [AUTOSYNC] CreateFile is not setting isFile flag in OmKeyInfo
- CDPD-59286: [AUTOSYNC] Reduce time of compaction pause during bootstrapping
- CDPD-59157: [ozone-cert-rotation] cert clean is unable to cleanup certificates LOCK error
- CDPD-58638: [AUTOSYNC] Ratis crash if a lot of directories deleted at once
- CDPD-58116: [AUTOSYNC] Ozone is supporting unicode volume and bucket names, potentially unintentionally
- CDPD-58047: Backport HIVE-23726 to CDP branches
- CDPD-57788: [AUTOSYNC] Snapdiff should read only keys with the bucket prefix
- CDPD-54981: [AUTOSYNC] [FSO] S3A compatibility - dfs -mkdir creates a zero byte file instead of a directory
- CDPD-52277: OM shutdown when creating key with malformed characters
- CDPD-52135: Error message is confusing when client fails to upload a key
- CDPD-51815: Fix the regex for key name validation
- CDPD-51329: Optimize block write path performance by reducing no of watchForCommit calls
- CDPD-48162: Getting exception for wildcard (*) search for database and table name
- CDPD-47138: distcp on OFS path failing with ClassNotFoundException when build is created using upstream ozone

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

- CVE-2023-39196
- CVE-2023-31582
- CVE-2020-28458
- CVE-2021-23445
- CVE-2023-34054
- CVE-2023-34062
- CVE-2023-46749
- CVE-2023-46750
- CVE-2023-41329
- CVE-2023-41327

Known issues in 7.1.9 CHF 3

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.1.9 CHF 3.

After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and 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, `ClassNotFoundException` for `org.cloudera.log4j.redactor.RedactorAppender` class was identified. The error message is harmless as the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(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 MPU. The problem is that 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.

OPSAPS-69846: If Ozone is installed with custom kerberos principals for its roles, operations on encrypted buckets can fail as Ranger KMS does not have its proxy users and groups configured for the custom S3 Gateway user.

Add the following configurations in Ranger-kms safety valve based on the custom s3g user. In this case , the user is s3gfoo0. The parameters are `hadoop.kms.proxyuser.s3gfoo0.hosts = *` `hadoop.kms.proxyuser.s3gfoo0.groups = *`

CDPD-66508: Shallow listing is enabled by default in 7.1.9. There is a bug in shallow listing that causes the below error when listing an empty directory in a LEGACY/OBS bucket:

error when listing an empty directory in a LEGACY/OBS bucket: `mkdir: getFileStatus on s3a:/testbucket/data/test: com.amazonaws.services.s3.model.AmazonS3Exception: Server Error`

(Service: Amazon S3; Status Code: 500; Error Code: 500 Server Error; Request ID: null; S3 Extended Request ID: null; Proxy: null), S3 Extended Request ID: null:500 Server Error: Server Error (Service: Amazon S3; Status Code: 500; Error Code: 500 Server Error; Request ID: null; S3 Extended Request ID: null; Proxy: null)

In S3 gateway log: Caused by: java.lang.IndexOutOfBoundsException: Index 0 out of bounds for length 0 at java.base/jdk.internal.util.Preconditions.outOfBounds(Preconditions.java:64) at java.base/jdk.internal.util.Preconditions.outOfBoundsCheckIndex(Preconditions.java:70) at java.base/jdk.internal.util.Preconditions.checkIndex(Preconditions.java:248) at java.base/java.util.Objects.checkIndex(Objects.java:372) at java.base/java.util.ArrayList.remove(ArrayList.java:535) at org.apache.hadoop.ozone.client.OzoneBucket\$KeyIterator.getNextShallowListOfKeys(OzoneBucket.java:1234) at org.apache.hadoop.ozone.client.OzoneBucket\$KeyIterator.getNextListOfKeys(OzoneBucket.java:1136) at org.apache.hadoop.ozone.client.OzoneBucket\$KeyIterator.hasNext(OzoneBucket.java:1110) at org.apache.hadoop.ozone.s3.endpoint.BucketEndpoint.get(BucketEndpoint.java:208) at jdk.internal.reflect.GeneratedMethodAccessor90.invoke(Unknown Source)

Disable shallow listing.

1. Log in to Cloudera Manager
2. Navigate to Clusters
3. Select the Ozone service
4. Go to Configurations
5. In S3 Gateway Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozone-site.xml, set ozone.s3g.list-keys.shallow.enabled = false.

CDPD-65801: This is an intermittent issue in native RocksDB tool which causes corruption to in-memory RocksDB metadata.

Set ozone.om.snapshot.load.native.lib to false and restart the OM.

CDPD-66142: When Solr is slow/down, Solr takes lot of time to respond to Recon Heatmap query or sometimes doesn't respond at all which makes Recon heatmap trying to load the heatmap data forever. This issue will be taken up in future releases and solution could be to introduce a health check for Solr or timeout the Recon query to Solr and show a meaningful message over Recon UI -> "Solr is not responding"

1. Stop Recon
2. Restart Solr
3. Start Recon

CDPD-66247: TestOzoneFileSystem.testListStatusOnKeyNameContainDelimiter is intermittent

None

CDPD-66261: When we have OBS or LEGACY bucket having keyPrefix starting with / like /readPath/ and fsPath configuration (ozone.om.enable.filesystem.paths) is enabled (true), then code flow will hit normalization of key during org.apache.hadoop.ozone.om.KeyManagerImpl#listKeys API call flow and normalized key will be vol/buck/readPath, but in keyTable, key will be saved as vol/buck//readPath/, so it does not match and listKeys API would not be able to retrieve the key with normalized key path. Fix for next CHF release : Check if keyPrefix starting with /, then while normalizing, do not remove the / slash at the beginning.

Can set the disable the fsPath configuration (ozone.om.enable.filesystem.paths) by setting as false in ozone-site.xml before running the test, but this may impact other test also if we set it at global level.

CDPD-66262: As HADOOP-16226 has not landed under CDH/hadoop 7.1.9 CHF3, trailing slashes in a string are not removed during keyName normalization. Consequently, the expected result of the test (TestObjectStoreWithFSO.testListKeysAtDifferentLevels) for listing keys with unnormalized keyNames in an FSO bucket does not match.

None

CDPD-66382: When Bucket layout is LEGACY and ozone.om.enable.filesystem.paths property is set to true, then delete will not work completely if keyName contains "/".

None.

CDPD-66252: du space calculation support for OBS and LEGACY (fsPath disabled).

du space for OBS buckets and LEGACY(fspath disabled) can be seen using CLI command.

OPSAPS-69539: CDP Runtime 7.1.9 does not support Oracle JDK 8u401 or OpenJDK 1.8.0_402 (8u402). Some services will fail to start. This can be a problem on RHEL 9.x as version 8u402 is the default OpenJDK 8 installed by the OS.

Workaround is to install an earlier version of JDK 8. For example Oracle jdk-8u291 / 1.8.0_291, or OpenJDK 8u292 / 1.8.0_292.

A fresh install of 7.1.9 or its CHF's does not allow user to bypass the Setup Database screen for YARN Queue Manager

YARN Queue Manager in Cloudera Data Platform (CDP) Private Cloud Base 7.1.9 CHF 2 does not require you to install a PostGres database, therefore users should not see the Setup Database screen and should be able to skip the Setup Database screen. With this known issue, users who are conducting a fresh install of 7.1.9 or its CHF's are not able to bypass the Setup Database screen as expected.

1. When conducting a fresh install of YARN Queue Manager in 7.1.9 or its CHF's, you must ensure that you have both CDP and Cloudera Manager upgraded to 7.1.9 CHF 2.
2. When you reach the Setup Database screen in the Cloudera Manager installation wizard for Queue Manager, enter any dummy values for the following fields:
 - a. Database name: configstore
 - b. Database Username: dbuser
 - c. Database Password: dbpassword
 - d. Database Hostname: localhost

YARN Queue Manager will not connect to PostGres with the above details and will fall back to the embedded database.

3. Run the following script command in a browser console to enable the Continue button:

```
document.querySelector('.btn.next').removeAttribute('disabled');
```

4. Click Continue and proceed with the YARN Queue Manager installation.
5. After installation is complete, SSH into the host that has Queue Manager installed, and run this command: `sed -i 's/migrationCompleted=true/migrationCompleted=false/' /var/lib/ha/hoop-yarn/migration.properties`



Note: Enable Queue Manager in the YARN configurations, and restart YARN.

6. Restart YARN Queue Manager.

CDPD-61524: Ozone Storage Container Manager fails to start on upgrading from CDP Private Cloud Base 7.1.6 to 7.1.9 CHF1. Also, if you have upgraded from CDP Private Cloud Base 7.1.6 to 7.1.7 or 7.1.8 and then to 7.1.9, the upgrade fails.

None. Cloudera recommends you to reach out to the Support before performing the upgrade to CDP Private Cloud Base 7.1.9.

CDPD-62254: Ozone is not supported on SLES15 with CHF1.

If your cluster has Ozone, Cloudera recommends you to not upgrade to 7.1.9 CHF1.

QAINFRA-18371: Conflict while installing libmysqlclient-devel on SLES 15

You may see an error such as the following while installing the mysql-devel and libmysqlclient-devel packages for setting up MariaDB as a backend database on SLES 15: File /usr/bin/mariadb_config from install of MariaDB-devel-<VERSION>.x86_64 conflicts with file from install of libmariadb-devel-3.1.21-150000.3.33.3.x86_64 (SLES Module Server Applications Updates)

While installing the mysql-devel and libmysqlclient-devel packages on SLES15, use the “--replac efiles” zypper switch or manually enter “yes” on the interactive pop-up that you see when the files are being overwritten.

CDPD-62834: Status of the deleted table is seen as ACTIVE in Atlas after the completion of navigator2atlas migration process

The status of the deleted table displays as ACTIVE.

None

CDPD-62837: During the navigator2atlas process, the hive_storagedesc is incomplete in Atlas

For the hive_storagedesc entity, some of the attributes are not getting populated.

None

CDPD-63690: RuntimeException encountered when generating snapshotDiff report between 2 snapshots

When snapshot feature is enabled, KeyDeletingService, SSTFilteringService and SnapDiff thread fall into a deadlock when accessing Snapshot Cache.

Restart the Ozone Manager.

CDPD-64238: Snapshot diff request failing when setting ozone.om.snapshot.db.max.open.files=-1

When snapshot feature is enabled, KeyDeletingService, SSTFilteringService and SnapDiff thread fall into a deadlock when accessing Snapshot Cache.

Restart the Ozone Manager.

OPSAPS-69481: Some Kafka Connect metrics missing from Cloudera Manager due to conflicting definitions

The metric definitions for kafka_connect_connector_task_metrics_batch_size_avg and kafka_connect_connector_task_metrics_batch_size_max in recent Kafka CSDs conflict with previous definitions in other CSDs. This prevents Cloudera Manager from registering these metrics. It also results in SMM returning an error. The metrics also cannot be monitored in Cloudera Manager chart builder or queried using the Cloudera Manager API.

Contact Cloudera support for a workaround.

CDPD-66938: [Analyze] [Atlas] test_time_range tests fail

When the Apache Atlas server is running on a node which has time zone other than UTC, there might be a time of day when the search results might differ if the relative CreateTime date range filters of TODAY, YESTERDAY, etc. are used. 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.

Use explicit date range filters instead of using relative date range filters, such as, TODAY, YESTERDAY.

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

CDPD-79160: NPE while deleting BusinessMetadata

If business metadata is created without adding any applicable types, a NullPointerException is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

CDPD-76789: Creating tag with name description throws java.lang.ClassCastException

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

UnsupportedClassVersionError

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.

CDPD-79099: Unable to import hive tables using import script in Public Cloud

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.

Technical Service Bulletins**TSB 2024-749: Possible Ozone Snapshot Chain Corruption in 7.1.9.0**

The Apache Ozone (Ozone) snapshots feature provides the ability to save consistent and immutable copies of the Ozone namespace. A chain of snapshots is tracked and maintained in the Ozone RocksDB that is utilized to calculate the difference between two snapshots required for the snapDiff API. A race condition on the snapshot delete path was found during testing. This condition might result in two consecutive snapshots pointing to the same parent, which breaks the snapshot chain.

Upstream JIRA

[HDDS-10524](#), [HDDS-10590](#), and [HDDS-9198](#)

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2024-749: Possible Ozone Snapshot Chain Corruption in 7.1.9.0](#)

TSB 2025-820: Potential Data Integrity Issues Found in Ozone

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

Knowledge article

For the latest update on this issue see the corresponding Knowledge 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

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.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

Cumulative hotfix 4

You can review the list of cumulative hotfixes that were shipped for CDP Private Cloud Base version 7.1.9 CHF4.

Cloudera Runtime 7.1.9.6 (Cumulative Hotfix 4) download URL:

Parcel Repository Location

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.6/parcels/
```

Fixed issues in 7.1.9 CHF 4

Know more about the cumulative hotfixes 4 for 7.1.9. This cumulative hotfix was released on March 11, 2024.



Note: Contact Cloudera Support for questions related to any specific hotfixes.

Following are the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p6.51045883

- CDPD-59217 - Upgraded Janino to 3.1.10
- KT-7527: Keytrustee parcel bits not available in unified code branch
- KT-7524: [FIPS+JDK11] Keytrustee should encrypt keys with "openssl pkcs8" command on RHEL 8.8
- KT-7523: Keytrustee is using non-centralized jackson version
- KT-7522: CLONE - Keytrustee HSM - Upgrade guava to the centralized CDPD version
- KT-7517: CLONE - Use the centralized logback version

- KT-7514: Postgres9 backup file gets overwritten in 2step upgrade
- COMPX-15962: Backport YARN-11369 Commons.compress throws an IllegalArgumentException with large uids after 1.21
- COMPX-15948: TestContinuousScheduling#testFairSchedulerContinuousSchedulingInitTime and TestFairScheduler#testNormalizationUsingQueueMaximumAllocation fails intermittently
- COMPX-15894: Backport MAPREDUCE-7468 Change add-opens flag's default value from true to false
- COMPX-15324: RM crashes if app is submitted to auto created queue with empty shortname
- COMPX-12937: QueueManager does not allow copying and pasting values while configuring queue sizes
- COMPX-6274: Fix failing unit test:
org.apache.hadoop.yarn.server.timelineservice.security.TestTimelineAuthFilterForV2.testPutTimelineEntities
- CDPD-66843: [7.1.9 CHF4 CLONE] - Provide an option to bypass evaluation of chained plugin if the parent plugin has applicable policy
- CDPD-66842: Ranger Admin server gives empty response when user with user-role tries to update lastname or email address
- CDPD-66798: [7.1.9 CHF4] Skip showing 'Page not found' for wrong value is provided to a api parameter in Login Session Tab
- CDPD-66796: [7.1.9 CHF4] Skip showing 'Page not found' page for INVALID_INPUT_DATA validation in User Profile
- CDPD-66790: Upgrade Jackson version to at least 2.15.0
- CDPD-66789: Centrilize and upgrade avro to 1.11.3 in streaming
- CDPD-66784: [7.1.9 CHF4] Update the execution of setServiceDef call in App.jsx
- CDPD-66782: [7.1.9 CHF4] Updating the "Something went wrong" page in Ranger React UI
- CDPD-66781: [7.1.9 CHF4] Audit logs for Masking policy is missing data mask type entry
- CDPD-66734: Backport ZOOKEEPER-4236 to 7.1.9 CHF4
- CDPD-66730: Phoenix-thirdparty - Upgrade Guava to 32.0.1 due to CVE-2023-2976
- CDPD-66630: spark build failure sles12
- CDPD-66604: HIVE-26961: Fixed improper replication metric count when the hive.repl.filter.transactions property is set to "true".
- CDPD-66525: hadoop: Upgrade logredactor to 2.0.16 (CDP 7.1.9)
- CDPD-66452: Enable unit-tests for Ranger gerrit PRs and canary
- CDPD-66432: HBase-Solr - Upgrade snakeyaml due to CVE-2022-1471
- CDPD-66424: [7.1.9] Upgrade Dropwizard to 2.1.11
- CDPD-66407: [7.1.7 SP3] Zeppelin is using non-centralized jackson version
- CDPD-66289: CLONE - 7.1.8x and 7.1.9.x CHF- Keytrustee-keyhsm - Upgrade Jetty to 9.4.53/10.0.17/11.0.17 due to CVE-2023-40167, CVE-2023-36479, CVE-2023-41900, CVE-2023-36478 and CVE-2023-44487
- CDPD-66279: Fixed an issue with Spark 3.3, that caused metastore connection to drop under certain circumstances. The drop was not affecting behavior, as it reconnects successfully, but caused noise in the logs and unplanned reconnects.
- CDPD-66262: [ERROR] org.apache.hadoop.ozone.om.TestObjectStoreWithFSO.testListKeysAtDifferentLevels
- CDPD-66261: [ERROR]
org.apache.hadoop.ozone.freon.TestOmBucketReadWriteKeyOps.testOmBucketReadWriteKeyOps
- CDPD-66257: [ERROR]
org.apache.hadoop.hdds.scm.container.metrics.TestSCMContainerManagerMetrics.testContainerOpsMetrics
- CDPD-66162: Phoenix Connectors - Upgrade Guava to 32.0.1 due to CVE-2023-2976
- CDPD-66156: Access Audits - Resource policy version used for of mask policy leading to Error page
- CDPD-66146: [7.1.9 CHF4] [Ranger React UI] Checkbox selection issue when clicking on permission label in tag-based permissions policy
- CDPD-66121: Mass deletion of 27TB data and the space is not reclaimed fully in ozone storage
- CDPD-66092: Fix Ranger Javapatch failure even if service-defs do not exist in ranger DB
- CDPD-65918: [7.1.9 CHF4] Inconsistent resource lookup behaviour with newly created service
- CDPD-65876: [7.1.9] Upgrade jackson version in SRM to 2.15.0

- CDPD-65875: [7.1.9] Upgrade Snakeyaml version in SRM to 2.0
- CDPD-65874: [7.1.9] Use centralized snappy-java version in Kafka
- CDPD-65870: [7.1.9] Upgrade Jackson version in Kafka
- CDPD-65841: [718, 719] Backport aarch64 related commits
- CDPD-65838: [7.1.9] Upgrade Jackson version in Cruise Control
- CDPD-65802: Kafka password is in clear text in application.properties backport
- CDPD-65800: Upgrade Sonarqube version after Gradle upgrade
- CDPD-65720: [AUTOSYNC] Remove io.dropwizard.metrics:metrics-ganglia dependency
- CDPD-65665: [7.1.9] Centralize streaming versions with common naming pattern
- CDPD-65634: [7.1.9] Upgrade Gradle to 8.5
- CDPD-65623: [7.1.9 CHF4] [Ranger React UI] Add inline assertions for displayName length in service creation / update form
- CDPD-65616: Not able to access zeppelin ui through Knox
- CDPD-65591: Iceberg replication is not working in 7.1.9 CHF3 stack
- CDPD-65590: IMPALA-12670: getIfPresent must throw the cause of error
- CDPD-65589: IMPALA-11501 Add flag to allow metadata-cache operations on masked tables
- CDPD-65586: [7.1.x] exclude log4j dependencies from spark-atlas-connector assembly
- CDPD-65583: [Spark] Backport CDPD-64232 to 7.1.7 SP2, 7.1.7 SP3, 7.1.8 and 7.1.9
- CDPD-65579: Avoid double XML escaping in SimpleDescriptorHandler
- CDPD-65458: Upgrade Gradle to 8
- CDPD-65433: Execute and read permissions granted to a user in different HDFS policies does not take effect.
- CDPD-65425: Upgrade Dropwizard version in SRM to 2.1.11
- CDPD-65402: Backport CDPD-64950 to 7.1.7.SP2 and 7.1.7.SP3
- CDPD-65293: [7.1.9] Upgrade Apache Ivy to 2.5.2 due to CVE-2022-46751
- CDPD-65239: Add missing libs in external_versions and centralize the same in zeppelin
- CDPD-65213: [AUTOSYNC] ManagedSecretKey.macInstances should not be ThreadLocal
- CDPD-65082: IMPALA-12584: Added backend configuration to restrict data file locations for Iceberg tables. The flag is enabled by default and Impala raises an error for Iceberg tables that consist of data files outside of the table directory.
- CDPD-65080: [7.1.9 CHF4] - Policy listing page experiences an unexpected reset to Access tab when attempting to filter the service and zone dropdown options
- CDPD-65079: [7.1.9 CHF4] - Optimize policy listing loader after session timeout and Audit Admin session ID modal loader
- CDPD-65077: [7.1.9 CHF4] - Optimize "plugins/definitions" API Call for Initial Load in Multiple Ranger-React Modules
- CDPD-64860: Upgrade Snakeyaml version in SRM to 2.0
- CDPD-64855: Upgrade Jackson version in SRM to 2.15.0
- CDPD-64803: [7.1.9 CHF4] - API calls for zones and services on initial landing in ZoneListing page is being called twice
- CDPD-64475: CDPD - Upgrade logback to 1.2.13/1.3.14/1.4.14 due to CVE-2023-6378 and CVE-2023-6481
- CDPD-64358: [AUTOSYNC] Pipeline.nodesInOrder should not be ThreadLocal
- CDPD-64235: CDPD - Upgrade Apache Derby to 10.14.3.0-cloudera1 due to CVE-2022-46337
- CDPD-63982: On 7.1.9 CHF FIPS+JDK11 cluster, Zeppelin service is not starting UP.
- CDPD-63464: [AUTOSYNC] EC: When Coordinator DN doing reconstruction, restart of target DN can lead to SCM crash
- CDPD-63463: [AUTOSYNC] EC: Recovering container cleanup at DN start is not happening due to NPE.
- CDPD-62890: [AUTOSYNC] Race condition in RocksDatabase
- CDPD-62717: [7.1.9 CHF4] Need to show Tag Policies for user when it has permission in "Tag Based Policies" module
- CDPD-62583: HMS Upgrade to 7.1.8.x or higher version fails if Hive log level is WARN
- CDPD-61475: Hadoop - Remove json-io due to CVE-2023-34610

- CDPD-60977: Hive - Upgrade Apache Ivy to 2.5.2 due to CVE-2022-46751
- CDPD-60950: [7.1.9 CHF4] - Error page 'Go back' button not redirecting to the right page
- CDPD-60830: HBase-Thirdparty - Upgrade Guava to 32.0.1 due to CVE-2023-2976
- CDPD-60742: open_connections and open_operations metrics not populated after hive service restart
- CDPD-60030: Hue : Stored Cross-Site Scripting in file name field
- CDPD-58770: Security - The config API endpoint returns the keyStorePassword
- CDPD-58580: CDPD - Upgrade Guava to 32.0.1 due to CVE-2023-2976
- CDPD-53885: Backport of HIVE-23444 - fixed in upstream hive and merged to cdw-master and cdpd-master branches. Backported to 717 SP3, 7.1.8 and 7.1.9 versions
- CDPD-53379: Grant permission in Impala engine not working with {owner} in ranger policy
- CDPD-50493: Sample Data from Table Browser in Hue launches expensive queries from the Impala Views
- CDPD-50047: Upgrade Schema Registry project to use Gradle 8
- CDPD-13292: externalize more common dependencies from Search, Solr, and Hbase-Indexer
- CDPD-11827: Backport ORC-616 "In Patched Base encoding, the value of headerThirdByte goes beyond the range of byte"
- OPSAPS-69481: Some Kafka Connect metrics missing from Cloudera Manager due to conflicting definitions
- OPSAPS-69458: Custom properties atlas.jaas.KafkaClient.option.password is in clear text in CDP cluster services

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

- CVE-2023-43642

Known issues in 7.1.9 CHF 4

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.1.9 CHF 4.

After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and 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, ClassNotFoundException for org.cloudera.log4j.redactor.RedactorAppender class was identified. The error message is harmless as the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

When using S3A committer fs.s3a.committer.name=directory with fs.s3a.committer.staging.conflict-mode=replace to write to FSO buckets, the client fails with the following error.

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(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 MPU. The problem is that 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.

OPSAPS-69846: If Ozone is installed with custom kerberos principals for its roles, operations on encrypted buckets can fail as Ranger KMS does not have its proxy users and groups configured for the custom S3 Gateway user.

Add the following configurations in Ranger-kms safety valve based on the custom s3g user. In this case, the user is s3gfoo0. The parameters are `hadoop.kms.proxyuser.s3gfoo0.hosts = *` `hadoop.kms.proxyuser.s3gfoo0.groups = *`

CDPD-66508: Shallow listing is enabled by default in 7.1.9. There is a bug in shallow listing that causes the below error when listing an empty directory in a LEGACY/OBS bucket:

```
Error when listing an empty directory in a LEGACY/OBS bucket:      mkdir: getFileStatus on s3
a://testbucket/data/test:      com.amazonaws.services.s3.model.AmazonS3Exception: Server Error
(Service: Amazon S3; Status Code: 500; Error Code: 500 Server Error; Request ID: null;
S3 Extended Request ID: null; Proxy: null), S3      Extended Request ID: null:500 Server Error:
Server Error (Service:      Amazon S3; Status Code: 500; Error Code: 500 Server Error; Request
ID: null; S3 Extended Request ID: null; Proxy: null)
```

```
In S3 gateway log: Caused by:      java.lang.IndexOutOfBoundsException: Index 0 out of bound
s for      length 0 at      java.base/jdk.internal.util.Preconditions.outOfBounds(Preconditions.ja
va:64)      at      java.base/jdk.internal.util.Preconditions.outOfBoundsCheckIndex(Preconditions.
java:70)      at      java.base/jdk.internal.util.Preconditions.checkIndex(Preconditions.java:248)
      at java.base/java.util.Objects.checkIndex(Objects.java:372) at      java.base/java.util.ArrayLis
t.remove(ArrayList.java:535) at      org.apache.hadoop.ozone.client.OzoneBucket$KeyIterator.g
etNextShallowListOfKeys(OzoneBucket.java:1234)      at      org.apache.hadoop.ozone.client
.OzoneBucket$KeyIterator.getNextListOfKeys(OzoneBucket.java:1136)      at      org.apache.
hadoop.ozone.client.OzoneBucket$KeyIterator.hasNext(OzoneBucket.java:1110)      at      or
g.apache.hadoop.ozone.s3.endpoint.BucketEndpoint.get(BucketEndpoint.java:208)      at      jdk.int
ernal.reflect.GeneratedMethodAccessor90.invoke(Unknown      Source)
```

Disable shallow listing by performing the following steps:

1. Log in to Cloudera Manager
2. Navigate to Clusters
3. Select the Ozone service
4. Go to Configurations
5. In S3 Gateway Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozone-site.xml, set `ozone.s3g.list-keys.shallow.enabled = false`.

CDPD-64394: OzoneManager may fail to start and have logs `java.lang.IllegalArgumentException: Trying to set updateID to XXX which is not greater than the current value of XXX for OMKeyInfo{XXX}` in OzoneManager log file.

Currently, this issue cannot be auto recovered. You must contact Cloudera support.

CDPD-65801: This is an intermittent issue in native RocksDB tool which causes corruption to in-memory RocksDB metadata.

Set `ozone.om.snapshot.load.native.lib` to false and restart the OM.

CDPD-66142: When Solr is slow/down, Solr takes lot of time to respond to Recon Heatmap query or sometimes doesn't respond at all which makes Recon heatmap trying to load the heatmap data forever. This issue will be taken up in future releases and solution could be to introduce a health check for Solr or timeout the Recon query to Solr and show a meaningful message over Recon UI -> "Solr is not responding"

1. Stop Recon
2. Restart Solr
3. Start Recon

CDPD-66247: TestOzoneFileSystem.testListStatusOnKeyNameContainDelimiter is intermittent

None

CDPD-66382: When Bucket layout is LEGACY and ozone.om.enable.filesystem.paths property is set to true, then delete will not work completely if keyName contains "/"

None.

CDPD-66252: du space calculation support for OBS and LEGACY (fsPath disabled).

du space for OBS buckets and LEGACY(fspath disabled) can be seen using CLI command.

OPSAPS-69539: CDP Runtime 7.1.9 does not support Oracle JDK 8u401 or OpenJDK 1.8.0_402 (8u402). Some services will fail to start. This can be a problem on RHEL 9.x as version 8u402 is the default OpenJDK 8 installed by the OS.

Workaround is to install an earlier version of JDK 8. For example Oracle jdk-8u291 / 1.8.0_291, or OpenJDK 8u292 / 1.8.0_292.

A fresh install of 7.1.9 or its CHF's does not allow user to bypass the Setup Database screen for YARN Queue Manager

YARN Queue Manager in Cloudera Data Platform (CDP) Private Cloud Base 7.1.9 CHF 2 does not require you to install a PostGres database, therefore users should not see the Setup Database screen and should be able to skip the Setup Database screen. With this known issue, users who are conducting a fresh install of 7.1.9 or its CHF's are not able to bypass the Setup Database screen as expected.

1. When conducting a fresh install of YARN Queue Manager in 7.1.9 or its CHF's, you must ensure that you have both CDP and Cloudera Manager upgraded to 7.1.9 CHF 2.
2. When you reach the Setup Database screen in the Cloudera Manager installation wizard for Queue Manager, enter any dummy values for the following fields:
 - a. Database name: configstore
 - b. Database Username: dbuser
 - c. Database Password: dbpassword
 - d. Database Hostname: localhost

YARN Queue Manager will not connect to PostGres with the above details and will fall back to the embedded database.

3. Run the following script command in a browser console to enable the Continue button:

```
document.querySelector('.btn.next').removeAttribute('disabled');
```

4. Click Continue and proceed with the YARN Queue Manager installation.
5. After installation is complete, SSH into the host that has Queue Manager installed, and run this command: `sed -i 's/migrationCompleted=true/migrationCompleted=false/' /var/lib/ha/ha-doop-yarn/migration.properties`



Note: Enable Queue Manager in the YARN configurations, and restart YARN.

6. Restart YARN Queue Manager.

CDPD-61524: Ozone Storage Container Manager fails to start on upgrading from CDP Private Cloud Base 7.1.6 to 7.1.9 CHF1. Also, if you have upgraded from CDP Private Cloud Base 7.1.6 to 7.1.7 or 7.1.8 and then to 7.1.9, the upgrade fails.

None. Cloudera recommends you to reach out to the Support before performing the upgrade to CDP Private Cloud Base 7.1.9.

CDPD-62254: Ozone is not supported on SLES15 with CHF1.

If your cluster has Ozone, Cloudera recommends you to not upgrade to 7.1.9 CHF1.

CDPD-62834: Status of the deleted table is seen as ACTIVE in Atlas after the completion of navigator2atlas migration process

The status of the deleted table displays as ACTIVE.

None

CDPD-62837: During the navigator2atlas process, the hive_storagedesc is incomplete in Atlas

For the hive_storagedesc entity, some of the attributes are not getting populated.

None

CDPD-63690: RuntimeException encountered when generating snapshotDiff report between 2 snapshots

When snapshot feature is enabled, KeyDeletingService, SSTFilteringService and SnapDiff thread fall into a deadlock when accessing Snapshot Cache.

Restart the Ozone Manager.

CDPD-64238: Snapshot diff request failing when setting ozone.om.snapshot.db.max.open.files=-1

When snapshot feature is enabled, KeyDeletingService, SSTFilteringService and SnapDiff thread fall into a deadlock when accessing Snapshot Cache.

Restart the Ozone Manager.

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

CDPD-79160: NPE while deleting BusinessMetadata

If business metadata is created without adding any applicable types, a NullPointerException is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

CDPD-76789: Creating tag with name description throws java.lang.ClassCastException

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

UnsupportedClassVersionError

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.

CDPD-79099: Unable to import hive tables using import script in Public Cloud

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.

Technical Service Bulletins

TSB 2024-749: Possible Ozone Snapshot Chain Corruption in 7.1.9.0

The Apache Ozone (Ozone) snapshots feature provides the ability to save consistent and immutable copies of the Ozone namespace. A chain of snapshots is tracked and maintained in the Ozone RocksDB that is utilized to calculate the difference between two snapshots required for the snapDiff API. A race condition on the snapshot delete path was found during testing. This condition might result in two consecutive snapshots pointing to the same parent, which breaks the snapshot chain.

Upstream JIRA

[HDDS-10524](#), [HDDS-10590](#), and [HDDS-9198](#)

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2024-749: Possible Ozone Snapshot Chain Corruption in 7.1.9.0](#)

TSB 2025-820: Potential Data Integrity Issues Found in Ozone

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

Knowledge article

For the latest update on this issue see the corresponding Knowledge 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

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.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

Cumulative hotfix 5

You can review the list of cumulative hotfixes that were shipped for CDP Private Cloud Base version 7.1.9 CHF5.

Cloudera Runtime 7.1.9.7 (Cumulative Hotfix 5) download URL:

Parcel Repository Location

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.7/parcels/
```

Fixed issues in 7.1.9 CHF 5

Know more about the cumulative hotfix 5 for 7.1.9. This cumulative hotfix was released on April 08, 2024.



Note: Contact Cloudera Support for questions related to hotfixes.

Following is the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p7.51778342

COMPX-16140: CDPD - Upgrade Spring Security to 5.7.11/5.8.7/6.0.7/6.1.4 due to CVE-2023-34042

Upgraded Spring Security to 5.7.11/5.8.7/6.0.7/6.1.4 due to CVE 2023-34042.

COMPX-11263: QM UI: Configuration modification shows old value temporarily

The configuration page displayed old values intermittently during an update. This issue is now fixed and only the new values entered by the user is displayed during the update.

CDPD-67558, YARN-11639: YARN RM stops assigning resources either because of ConcurrentModificationException or NPE in PriorityUtilizationQueueOrderingPolicy

When dynamic queue creation was enabled in weight mode and the deletion policy coincides with the `PriorityQueueResourcesForSorting`, YARN RM stopped assigning resources because of either `ConcurrentModificationException` or `NPE` in `PriorityUtilizationQueueOrderingPolicy`. This issue is now fixed.

CDPD-67507: Use Name validation regex instead of service name validation regex for Display name

The regex validator for validating the service display name in the service edit form did not allow a space character. This fix now allows a space character in the service display name.

CDPD-67433: IMPALA-12878 TestResultSpoolingCancellation.test_cancellation failed in UBSAN build

A rare scenario where a query is closed by the client then closed again (most clients prevent this happening), resulted in an error message Query not yet running. This fix restores the previous message Invalid or unknown query handle.

CDPD-67313: [7.1.x] Timezone value not updated in Livy

While serialising the Spark results from the Java object into JSON text, Livy was setting the timezone to UTC. The Livy service code is now fixed to use a custom timezone based on the configuration instead of always using UTC.

CDPD-67278: Backport KNOX-3012: Fix the DN links on the Ozone SCM UI

There was a change in Ozone, which caused the links of the DataNodes (DN) not to route through Knox on the Ozone SCM UI. With this fix the DN links redirect to the correct Knox URLs again.

CDPD-67225: Zeppelin - Upgrade Spring Framework to 6.1.4/6.0.17/5.3.32 due to CVE-2024-22243

Upgraded Spring Framework to 6.1.4/6.0.17/5.3.32 due to CVE-2024-22243.

CDPD-67220: [Regression] Oozie HTTPS notification fails if SSL is not set in Oozie

Oozie fails to assemble the notification request if the notification URL is secure (uses HTTPS) but no SSL is configured for Oozie server. This issue is now resolved.

CDPD-67193: The inactivityTimeout is reset when user updates the profile from UserProfile page

Fixed an issue of not resetting inactivityTimeout to a default value of 15 minutes when the user updated the profile on the **User Profile** page of the Ranger Admin UI.

CDPD-67023: [Ranger React UI] Audit UI improvements with respect to values overflowing into other columns

In the Ranger React UI, for certain columns in the audits pages, the value overflowed into the next columns if the text length was long. This issue is now fixed, and the following columns in the specified audit pages are modified to prevent the overflow into next column:

- Access Audits - Service Name and Cluster Name
- Plugin Status Audits - Service Name field
- Login sessions Audits - Login Id field

CDPD-66997: [AUTOSYNC] Recon - UnsupportedOperationException while merging Incremental Container Reports

A UnsupportedOperationException was displayed while merging incremental container reports. This issue is now fixed.

CDPD-66963, CDPD-66773: [AUTOSYNC] NPE causes OM crash in Snapshot Purge request

Ozone now ignores a purge request if there is a snapshot purge request for an already purged snapshot.

CDPD-66934: Display query information for Show databases/schemas command on Ranger Admin UI

In the 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 did not display. This issue is now fixed and the UI now displays the query / operation information for access audits even if the resource type is null.

CDPD-66927: HDFS authorization logic for directory hierarchy rooted at "/" is incorrect

There was an issue with the Ranger authorization logic for the HDFS commands that required authorization of the entire directory hierarchy rooted at the specified directory argument. This argument was incorrect due to incorrect computation of the sub-directory paths. The paths of sub-directories to be authorized contained an extra / character, leading to incorrect authorization result. This issue is now fixed.

CDPD-66917: [AUTOSYNC] Upgrade aws-java-sdk to 1.12.661

Upgrading aws-java-sdk to 1.12.661 version removes ion-java dependency from aws-java-sdk which caused CVE-2024-21634.

CDPD-66843: [7.1.9 CHF5 CLONE] - Provide an option to bypass evaluation of chained plugin if the parent plugin has applicable policy

When a chained plugin (such as Hive) is configured, every access request processed by the parent plugin (such as HDFS/Ozone/S3) is also processed by the chained plugin. This feature now supports

a configuration parameter `ranger.plugin.bypass.chained.plugin.evaluation.if.access.is.determined` with the default value as `FALSE`. When set to `TRUE`, the evaluation of the chained plugin is skipped when an applicable policy is found by the parent plugin. This issue is now fixed.

CDPD-66842: Ranger Admin server provides empty response when user with user-role tries to update lastname or email address

An error response with a message is now displayed when a user with a user-role tries to add or 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. Enabling performance logging for the module helps in measuring the amount of time spent during execution of various methods or functions during its operation. For achieving more precise time measurement, this feature supports nanosecond precision when the JVM version supports it.

CDPD-66798: [7.1.9 CHF5] Skip showing Page not found for wrong value is provided to a API parameter in Login Session Tab

There was an issue where Page Not Found was displayed when a user entered a text value to a search an API parameter **IP** in the **Login Sessions** under **Audits**. This issue is fixed and a server-side response is now displayed for invalid values as an alert on the **Login Sessions** tab.

CDPD-66796: [7.1.9 CHF5] Skip showing Page not found page for INVALID_INPUT_DATA validation in User Profile

A Page Not Found error message was displayed when a user provided invalid form values during profile update. This issue is now fixed and a server-side response is displayed as an alert on the **User Profile** window.

CDPD-66784: Update the execution of setServiceDef call in App.jsx

Removed unused code related to setServiceDef call in App.jsx.

CDPD-66782: Updating the Something went wrong page in Ranger React UI

The Something went wrong message was displayed when there was an error in the React JS code that was used to load Ranger Admin UI. This issue is now fixed.

CDPD-66781: Audit logs for Masking policy is missing data mask type entry

Fixed issue of showing Audit log for custom data mask type when added or updated into a policyItem of Masking policy.

CDPD-66725: Knox - Upgrade Okio to 3.4.0 due to CVE-2023-3635

Upgraded Okio to 3.4.0 due to CVE-2023-3635.

CDPD-66719: Ranger - Upgrade Spring Security to 5.7.11/5.8.7/6.0.7/6.1.4 due to CVE-2023-34042

Upgraded Spring Security to 5.7.11

CDPD-66568: Export/Import : changeMarker is not set to entity's lastupdate time or its closer timestamp value

When a Hive table entity was exported using a fetch type incremental with changeMarker 0, after exporting, the changeMarker in the export response was not set to a recent timestamp. This issue is now fixed, and the changeMarker is now set to a closer timestamp value during an export or import.

CDPD-66538: [AUTOSYNC] Metadata are not updated when keys are overwritten

There was an issue when an object was created with the same key name as one already present in the database. The request was forwarded to the Ozone Manager (OM) side of the code, specifically to the OmKeyRequest class, containing a method called prepareFileInfo(). This method persists the data to the openKeyTable. Initially, the method checked if a key with the same name exists. If exists, new data size, modification time, updateID, and replicationConfig was updated. However, the metadata of the overridden file was not updated. Consequently, the old metadata stored earlier is retained.

This issue is now fixed and the changes involve extracting new metadata from the KeyArgs object and comparing it with the existing metadata in the OmKeyInfo object. Any new or modified metadata entries are then updated in the OmKeyInfo object. Also, metadata entries not mentioned in the overwrite operation are retained, ensuring the preservation of existing metadata.

CDPD-66509: [7.1.9]BackPort to 7.1.9x branch

Upgraded common-dbc2 to 2.1.0 and commons-pool2 to 2.12.0

CDPD-66423: Backport HIVE-25986 to 7.1.9.x branches

The statement ID was incorrect if the table was an insert only ACID table and the LOAD IN PATH command was used to load the data. Because of this incorrect statement ID, the delta file path also contained a name, which was incompatible with other systems such as Impala. This issue is now fixed and the correct statement ID is now generated.

CDPD-66417: Upgrade Prometheus to 2.45.3 due to CVEs

Upgraded Prometheus to 2.45.3 to address CVE-2023-44487 and CVE-2023-45142.

CDPD-66358: HS2 logs having WARN logs from RangerHiveAuthorizer regarding connection to HMS for fetching hive object owner

This fix addresses the issue of HS2 logs having huge number of WARN logs.

CDPD-66243: [Knox] Invalid binary character logged in gateway.log

Upgraded the libpam4j dependency to fix a bug that resulted in group names with invalid characters.

CDPD-65969: A change in the message for ozone admin cert list subcommand count limits

Listing of certificates is performed in a batch with a default size of 20. A few certificates were not displayed if there were more than 20. This issue is now fixed and a warning message is displayed if the batch size is limiting the amount of certificates displayed and an option is provided to increase the batch size.

CDPD-65808: [7.1.9 CHF5 CLONE] - Performance degradation while retrieving mapped Hive resource for S3 location.

This fix improves the performance of RMS access evaluation while retrieving mapped Hive resource for Ozone locations (that is, Ozone keys).

CDPD-65616: Not able to access Zeppelin UI through Knox

Added JVM arguments to expose hidden internal classes required by the Ranger plug-in.

CDPD-65001: [AUTOSYNC] Pass TransactionInfo in OzoneManagerRequestHandler.handleWriteRequest

In OzoneManagerRequestHandler.handleWriteRequest, only transactionLogIndex was passed without the term. This issue is now fixed and the TransactionInfo (includes both term and index) is now passed. Thus, avoiding the recalculation of the term later on.

CDPD-64938: [AUTOSYNC] Remove RatisSnapshotInfo

Fixed the inconsistency with Ratis term and index when values are printed using the toString() command and when operations are running in parallel.

CDPD-64822: [AUTOSYNC] Move add response in doubleBuffer from validateAndUpdateCache to handleWriteRequest

Fixed an issue to ensure every response returned from validateAndUpdateCache is added to DoubleBuffer.

CDPD-64626: CLONE - Ranger - Upgrade aws-java-sdk-bundle to 1.12.599 due to CVE-2023-44487

Upgraded aws-java-sdk-bundle to 5.7.11

CDPD-64394: [AUTOSYNC] OzoneManagerStateMachine should put all failed write requests into OzoneManagerDoubleBuffer

Fixed the OzoneManager (OM) restart failure issue due to failed OM write request's response not added to OzoneManagerDoubleBuffer.

CDPD-64153: [AUTOSYNC] Tool to fix corrupted snapshot chain

This tool is a workaround to fix the snapshot chain corruption issue until the root cause is identified and fixed for the snapshot chain corruption.

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.

CDPD-63687: Deleted resource mapping is not removed from the plugin's cache

When a storage (HDFS/Ozone/S3) is configured to use RMS, the storage locations of Hive/Impala database/table objects are maintained by the RMS server, and provided to the Ranger authorizer running in the storage service. This feature ensures that when a Hive database/table is removed, mapping information for the removed object is cleared from the resource-mappings provided to the storage service.

CDPD-63039: IMPALA-12528 test_hdfs_scanner_thread_non_reserved_bytes may occasionally fail

Fixed unit test issue at `test_hdfs_scanner_thread_non_reserved_bytes`.

CDPD-60459: HueQP - Fixing NPE for adminUser in facets api

A `NullPointerException` was displayed in `facets/` API. This affected users of Hive job browser (using QP) with impact on all Hue/Hive users, and no impact to other components. This issue is now fixed.

CDPD-48298: CLONE - Knox - Upgrade Guava: Google Core Libraries for Java to v28.2/31.1-jre due to low CVEs

Updated the Guava dependency to get the fix for CVE-2020-8908.

CDPD-46225: Security Zone policies version increases by two when you update its resource.

The issue where updating a resource resulted in the Security Zone policies version incrementing by two is now resolved.

CDPD-44220: Livy - Missing deploy mode param at Spark submit

Fixed occasional issues with session recovery/HA failover on FIPS clusters.

CDPD-67307: Upgrade NodeJS version to 20.11.1 due to CVEs

NodeJS has been upgraded to version 20.11.1 to apply fixes to CVEs.

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

- CVE-2023-36478
- CVE-2023-26048
- CVE-2023-26049
- CVE-2023-40167
- CVE-2023-41900

Technical Service Bulletins**TSB 2024-752: Dangling delete issue in Spark `rewrite_data_files` procedure causes incorrect results for Iceberg V2 tables**

For the latest update on this issue see the corresponding Knowledge article: [TSB 2024-752: Dangling delete issue in Spark `rewrite\_data\_files` procedure causes incorrect results for Iceberg V2 tables](#).

Known issues in 7.1.9 CHF 5

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.1.9 CHF 5.

After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and 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, `ClassNotFoundException` for `org.cloudera.log4j.redactor.RedactorAppender` class was identified. The error message is harmless as the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(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 MPU. The problem is that 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.

CDPD-67095: In Knox on the Ozone SCM UI, the datanode links are not rewritten and they are not routing through Knox.

In Knox on the Ozone SCM UI the datanode links are not rewritten and they are not routing through Knox.

None.

OPSAPS-69846: If Ozone is installed with custom kerberos principals for its roles, operations on encrypted buckets can fail as Ranger KMS does not have its proxy users and groups configured for the custom S3 Gateway user.

Add the following configurations in Ranger-kms safety valve based on the custom s3g user.
In this case , the user is s3gfoo0. The parameters are `hadoop.kms.proxyuser.s3gfoo0.hosts = *`
`hadoop.kms.proxyuser.s3gfoo0.groups = *`

CDPD-66508: Shallow listing is enabled by default in 7.1.9. There is a bug in shallow listing that causes the below error when listing an empty directory in a LEGACY/OBS bucket:

Error when listing an empty directory in a LEGACY/OBS bucket: `mkdir: getFileStatus on s3a://testbucket/data/test: com.amazonaws.services.s3.model.AmazonS3Exception: Server Error (Service: Amazon S3; Status Code: 500; Error Code: 500 Server Error; Request ID: null; S3 Extended Request ID: null; Proxy: null), S3 Extended Request ID: null:500 Server Error: Server Error (Service: Amazon S3; Status Code: 500; Error Code: 500 Server Error; Request ID: null; S3 Extended Request ID: null; Proxy: null)`

In S3 gateway log: Caused by: `java.lang.IndexOutOfBoundsException: Index 0 out of bounds for length 0 at java.base/jdk.internal.util.Preconditions.outOfBounds(Preconditions.java:64) at java.base/jdk.internal.util.Preconditions.outOfBoundsCheckIndex(Preconditions.java:70) at java.base/jdk.internal.util.Preconditions.checkIndex(Preconditions.java:248) at java.base/java.util.Objects.checkIndex(Objects.java:372) at java.base/java.util.ArrayList.remove(ArrayList.java:535) at org.apache.hadoop.ozone.client.OzoneBucket$KeyIterator.getNextShallowListOfKeys(OzoneBucket.java:1234) at org.apache.hadoop.ozone.client.OzoneBucket$KeyIterator.getNextListOfKeys(OzoneBucket.java:1136) at org.apache.hadoop.ozone.client.OzoneBucket$KeyIterator.hasNext(OzoneBucket.java:1110) at org.apache.hadoop.ozone.s3.endpoint.BucketEndpoint.get(BucketEndpoint.java:208) at jdk.internal.reflect.GeneratedMethodAccessor90.invoke(Unknown Source)`

Disable shallow listing by performing the following steps:

1. Log in to Cloudera Manager
2. Navigate to Clusters
3. Select the Ozone service
4. Go to Configurations
5. In S3 Gateway Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozone-site.xml, set `ozone.s3g.list-keys.shallow.enabled = false`.

CDPD-66247: TestOzoneFileSystem.testListStatusOnKeyNameContainDelimiter is intermittent

None

CDPD-66252: du space calculation support for OBS and LEGACY (fsPath disabled).

`du space` for OBS buckets and LEGACY(fspath disabled) can be seen using CLI command.

OPSAPS-69539 : CDP Runtime 7.1.9 does not support Oracle JDK 8u401 or OpenJDK 1.8.0_402 (8u402). Some services will fail to start. This can be a problem on RHEL 9.x as version 8u402 is the default OpenJDK 8 installed by the OS.

Workaround is to install an earlier version of JDK 8. For example Oracle `jdk-8u291 / 1.8.0_291`, or OpenJDK `8u292 / 1.8.0_292`.

A fresh install of 7.1.9 or its CHF's does not allow user to bypass the Setup Database screen for YARN Queue Manager

YARN Queue Manager in Cloudera Data Platform (CDP) Private Cloud Base 7.1.9 CHF 2 does not require you to install a Postgres database, therefore users should not see the Setup Database screen and should be able to skip the Setup Database screen. With this known issue, users who are conducting a fresh install of 7.1.9 or its CHF's are not able to bypass the Setup Database screen as expected.

1. When conducting a fresh install of YARN Queue Manager in 7.1.9 or its CHF's, you must ensure that you have both CDP and Cloudera Manager upgraded to 7.1.9 CHF 2.

2. When you reach the Setup Database screen in the Cloudera Manager installation wizard for Queue Manager, enter any dummy values for the following fields:
 - a. Database name: configstore
 - b. Database Username: dbuser
 - c. Database Password: dbpassword
 - d. Database Hostname: localhost

YARN Queue Manager will not connect to PostGres with the above details and will fall back to the embedded database.

3. Run the following script command in a browser console to enable the Continue button:

```
document.querySelector('.btn.next').removeAttribute('disabled');
```

4. Click Continue and proceed with the YARN Queue Manager installation.
5. After installation is complete, SSH into the host that has Queue Manager installed, and run this command: `sed -i 's/migrationCompleted=true/migrationCompleted=false/' /var/lib/ha/hoop-yarn/migration.properties`



Note: Enable Queue Manager in the YARN configurations, and restart YARN.

6. Restart YARN Queue Manager.

CDPD-61524: Ozone Storage Container Manager fails to start on upgrading from CDP Private Cloud Base 7.1.6 to 7.1.9 CHF1. Also, if you have upgraded from CDP Private Cloud Base 7.1.6 to 7.1.7 or 7.1.8 and then to 7.1.9, the upgrade fails.

None. Cloudera recommends you to reach out to the Support before performing the upgrade to CDP Private Cloud Base 7.1.9.

CDPD-62254: Ozone is not supported on SLES15 with CHF1.

If your cluster has Ozone, Cloudera recommends you to not upgrade to 7.1.9 CHF1.

CDPD-62834: Status of the deleted table is seen as ACTIVE in Atlas after the completion of navigator2atlas migration process

The status of the deleted table displays as ACTIVE.

None

CDPD-62837: During the navigator2atlas process, the hive_storagedesc is incomplete in Atlas

For the hive_storagedesc entity, some of the attributes are not getting populated.

None

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet `atlas.authentication.method.file` is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

CDPD-79160: NPE while deleting BusinessMetadata

If business metadata is created without adding any applicable types, a `NullPointerException` is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

CDPD-66938: [Analyze] [Atlas] test_time_range tests fail

When the Apache Atlas server is running on a node which has time zone other than UTC, there might be a time of day when the search results might differ if the relative `CreateTime` date range filters of `TODAY`, `YESTERDAY`, etc. are used. 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.

Use explicit date range filters instead of using relative date range filters, such as, TODAY, YEST ERDAY.

CDPD-76789: Creating tag with name description throws java.lang.ClassCastException

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

UnsupportedClassVersionError

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.

CDPD-79099: Unable to import hive tables using import script in Public Cloud

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.

Technical Service Bulletins**TSB 2024-749: Possible Ozone Snapshot Chain Corruption in 7.1.9.0**

The Apache Ozone (Ozone) snapshots feature provides the ability to save consistent and immutable copies of the Ozone namespace. A chain of snapshots is tracked and maintained in the Ozone RocksDB that is utilized to calculate the difference between two snapshots required for the snapDiff API. A race condition on the snapshot delete path was found during testing. This condition might result in two consecutive snapshots pointing to the same parent, which breaks the snapshot chain.

Upstream JIRA

[HDDS-10524](#), [HDDS-10590](#), and [HDDS-9198](#)

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2024-749: Possible Ozone Snapshot Chain Corruption in 7.1.9.0](#)

TSB 2025-820: Potential Data Integrity Issues Found in Ozone

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

Knowledge article

For the latest update on this issue see the corresponding Knowledge 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

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.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

Cumulative hotfix 6

You can review the list of cumulative hotfixes that were shipped for CDP Private Cloud Base version 7.1.9 CHF6.

Cloudera Runtime 7.1.9.9 (Cumulative Hotfix 6) download URL:

Parcel Repository Location

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.9/parcels/
```

Fixed issues in 7.1.9 CHF 6

Know more about the cumulative hotfix 6 for 7.1.9. This cumulative hotfix was released on April 26, 2024.



Note: Contact Cloudera Support for questions related to hotfixes.

Following is the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p9.52289703

CDPD-60645: Add more logging for snapshot purge.

Extensive logging is added on the Snapshot purge code path for operational usage.

CDPD-67651: Recon UI changes for Solr Health Check.

Health Check in Recon Dashboard for Solar is Healthy or Unhealthy. This issue is now resolved.

COMPX-16221: CLONE [7.1.9 | 7.2.19]- Queue Manager - Upgrade protobuf-java to 3.16.3/3.19.6/3.20.3/3.21.7 due to CVE-2022-3171

Upgraded protobuf-java to 3.16.3/3.19.6/3.20.3/3.21.7 due to CVE-2022-3171.

CDPD-68597: Need to include HDDS-10430 in CHF6

A race condition resulted in failure of serialization of Pipeline information at Ozone Manager when using topology-aware read. This issue is now fixed.

CDPD-68541: SnakeYaml upgrade due to CVE in schema_registry_ranger_plugin

Upgraded the SnakeYaml version referenced in the schema_registry_ranger_plugin component, due to CVE-2022-1471.

CDPD-68466, CDPD-68129: Change snapshot purge request from batch to single snapshot at a time

In Ozone Snapshot, the previous snapshot was retained for each SnapshotInfo entry in the SnapshotInfoTable. On the snapshot create code path, the latest snapshot was added as the previous snapshot. During snapshot purging, there was race condition on the snapshot purge path, where two consecutive snapshots pointed to the same previous snapshot causing snapshot chain corruption. This issue is now fixed.

CDPD-68339: [AUTOSYNC] Ozone Recon - Filter EMPTY MISSING Containers in UnHealthy State Containers API

Recon UnHealthy State Containers API was not filtering the empty containers and the containers were reported as MISSING. This issue is now fixed.

CDPD-67945: [AUTOSYNC] Link rocksdb lib to Ozone rocksdb tools lib relative path instead of absolute path

The native library load failed because the rocksdb library was linked to the absolute path. This issue is now fixed by linking the rocksdb library to the Ozone rocksdb tools library relative path instead of the absolute path.

CDPD-67944: Make Rocksdb tools native lib compatible with all chipsets within the arch

While building rocksdb tools library, the makefile of rocksdb added the native optimization gcc flags when compiling the rocksdb code. This made the build code incompatible with other chipsets. This issue is now fixed and the rocksdb tools native library is now compatible with all chipsets within the arch.

CDPD-67864: Ranger - Upgrade Spring Security to 5.7.12/5.8.11/6.1.8/6.2.3 due to CVE-2024-22257

Upgraded Spring Security to 5.7.12 due to CVE-2024-22257.

CDPD-67828: Hadoop - Upgrade Nimbus-JOSE-JWT to 9.37.3 due to CVE-2023-52428

Upgraded Nimbus-JOSE-JWT to 9.37.3 due to CVE-2023-52428

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.

CDPD-67820: Backport HIVE-26955 to CDH-7.1.9.x

Select query failed when the decimal column data type was changed to string, char or varchar in Parquet format. This issue is now fixed.

CDPD-67819: Backport HIVE-26320 to CDH-7.1.9.x

A query involving case statement with two or more conditions lead to incorrect result for tables with parquet format. This issue is now fixed.

CDPD-67750: Ranger - Upgrade telemetry to 1.36.0

Upgraded Telemetry version to 1.36.0.

CDPD-67749: Ranger - Upgrade protobuf-java to 3.21.7 due to CVE-2022-3171

Upgraded protobuf-java to 3.21.7 due to CVE-2022-3171.

CDPD-67746: Ranger - Upgrade Nimbus-JOSE-JWT to 9.37.3 due to CVE-2023-52428

Upgrade Nimbus-JOSE-JWT to 9.31

CDPD-67744: Ranger - Exclude Apache Derby from ranger-rms module due to CVE-2022-46337

Exclude apache derby from ranger-rms module

CDPD-67440, CDPD-67727: [AUTOSYNC] Snapshot chain corruption

In Ozone Snapshot, we keep snapshot's previous snapshot for each SnapshotInfo entry in the SnapshotInfoTable. On the snapshot create code path, we use the latest snapshot added as the previous snapshot. And on snapshot purging, we update the next snapshot's previous to the snapshotToPurge's previous to maintain the chain. (Similar to deleting a node from a LinkedList/ DoublyLinkedList). There is a race condition on the snapshot purge path because of which sometimes, two consecutive snapshots point to the same previous snapshot and cause snapshot chain corruption.

CDPD-67593: [Spark Compaction] Count function in impala giving incorrect results for V2 Iceberg tables

The count function in Impala gave incorrect results for V2 Iceberg tables. This issue is now fixed by turning off the optimized count(*) feature.

CDPD-67070: Backport HBASE-27230 to 7.1.9 CHF5

As per the design, write-ahead-log synchronisation (WAL Sync) should always succeed or abort. There is no failure state. Hence, if the WAL sync runs into timeout exception, the region-server should abort. The timeout for the sync is generally very large (5 minutes). Hence, if the operation does not complete by that time and runs into timeout, then, the region server should abort.

CDPD-66926: OM startup failure : OMLeaderNotReadyException leader is not getting ready

Ozone health was in a bad state and all the requests failed with errors as OM leaders did process any request. This issue is now fixed.

CDPD-66893: Hadoop - Upgrade moment.js to 2.29.4 due to CVE-2022-24785, CVE-2022-31129

Upgraded moment.js to version 2.29.4 due to CVE-2022-24785 and CVE-2022-31129.

CDPD-66890: TestMapReduceIndexerTool localfs test error

On FIPS environments the localfs functionality of both MRIT and HBase MRIT was not working because the underlying JSCH library did not support SHA2. This issue is now fixed by replacing JSCH with native implementation.

CDPD-66743: [AUTOSYNC] Recon - Handle the pre-existing missing empty containers in clusters

There were EMPTY MISSING containers (zero keys), and those were reported as MISSING by Recon. This fixes and handles the pre-existing empty missing containers.

CDPD-66668: Kafka_connect_ext - Upgrade commons-compress to 1.26.0 due to CVE-2023-42503, CVE-2024-25710 and CVE-2024-26308

Upgrade Commons-Compress to version 1.26.0 due to CVE-2023-42503, CVE-2024-25710 and CVE-2024-26308.

CDPD-65188: [AUTOSYNC] Remove applyTransactionMap and ratisTransactionMap from OzoneManagerStateMachine

The applyTransactionMap and the ratisTransactionMap in OzoneManagerStateMachine is now replaced by two TermIndex values:

- lastApplied (implemented by BaseStateMachine) : keeps track of the last flushed TermIndex from OzoneManagerDoubleBuffer
- lastNotified(a new field in OzoneManagerStateMachine): keeps track the last term, index passed from notifyTermIndexUpdated(..)

CDPD-65178: [AUTOSYNC] OzoneManagerRatisServer.getServer() should return Division**CDPD-65158: [AUTOSYNC] NumKeys metric not decremented on FSO directory delete**

The NumKeys metric is now decremented for all the keys within a directory when the directory is deleted. This is done in the background as the directory is cleaned up after the initial deletion.

CDPD-64334: Ranger - Upgrade Bouncy Castle to 1.78 due to CVE-2023-33202, CVE-2023-33201, CVE-2024-29857, CVE-2024-30171 and CVE-2024-30172

Upgraded Bouncy Castle library version to 1.77 due to CVE-2023-33202, CVE-2023-33201, CVE-2024-29857, CVE-2024-30171 and CVE-2024-30172.

CDPD-63605: Verbose log messages regarding pipeline creation failure

New pipelines were not created because logs were flooded with messages regarding SCM and when the cluster had reached maximum pipelines per datanode configuration limit. This issue is now fixed by limiting such verbose messages.

CDPD-50862: Extend usage of fire_listener_event API to HS2/Spark to generate events on DML queries

The data written from Spark now generates an Insert event in the HMS notification log table. This is useful for external services (listeners) such as Impala to know the current condition of an external table.

CDPD-42610: Make RetryPolicy Configurable in Read Key path

Ozone client retries read operation three times when Datanodes are unavailable temporarily. The retry policy is now configurable using the client configurations `ozone.client.read.max.retries(3` by default) and `ozone.client.read.retry.interval(1` by default).

Common Vulnerabilities and Exposures (CVE) that is fixed in this CHF:

- CVE-2023-44483

Known issues in 7.1.9 CHF 6

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.1.9 CHF 6.

After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and 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, `ClassNotFoundException` for `org.cloudera.log4j.redactor.RedactorAppender` class was identified. The error message is harmless as the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

When using S3A committer `fs.s3a.committer.name=directory` with `fs.s3a.committer.staging.conflict-mode=replace` to write to FSO buckets, the client fails with the following error.

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(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 MPU. The problem is that 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.

CDPD-67095: In Knox on the Ozone SCM UI, the datanode links are not rewritten and they are not routing through Knox.

In Knox on the Ozone SCM UI the datanode links are not rewritten and they are not routing through Knox.

None.

CDPD-66508: Shallow listing is enabled by default in 7.1.9. There is a bug in shallow listing that causes the below error when listing an empty directory in a LEGACY/OBS bucket:

Error when listing an empty directory in a LEGACY/OBS bucket: mkdir: getFileStatus on s3a://testbucket/data/test: com.amazonaws.services.s3.model.AmazonS3Exception: Server Error (Service: Amazon S3; Status Code: 500; Error Code: 500 Server Error; Request ID: null; S3 Extended Request ID: null; Proxy: null), S3 Extended Request ID: null:500 Server Error: Server Error (Service: Amazon S3; Status Code: 500; Error Code: 500 Server Error; Request ID: null; S3 Extended Request ID: null; Proxy: null)

In S3 gateway log: Caused by: java.lang.IndexOutOfBoundsException: Index 0 out of bounds for length 0 at java.base/jdk.internal.util.Preconditions.outOfBounds(Preconditions.java:64) at java.base/jdk.internal.util.Preconditions.outOfBoundsCheckIndex(Preconditions.java:70) at java.base/jdk.internal.util.Preconditions.checkIndex(Preconditions.java:248) at java.base/java.util.Objects.checkIndex(Objects.java:372) at java.base/java.util.ArrayList.remove(ArrayList.java:535) at org.apache.hadoop.ozone.client.OzoneBucket\$KeyIterator.getNextShallowListOfKeys(OzoneBucket.java:1234) at org.apache.hadoop.ozone.client.OzoneBucket\$KeyIterator.getNextListOfKeys(OzoneBucket.java:1136) at org.apache.hadoop.ozone.client.OzoneBucket\$KeyIterator.hasNext(OzoneBucket.java:1110) at org.apache.hadoop.ozone.s3.endpoint.BucketEndpoint.get(BucketEndpoint.java:208) at jdk.internal.reflect.GeneratedMethodAccessor90.invoke(Unknown Source)

Disable shallow listing by performing the following steps:

1. Log in to Cloudera Manager
2. Navigate to Clusters
3. Select the Ozone service
4. Go to Configurations
5. In S3 Gateway Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozone-site.xml, set `ozone.s3g.list-keys.shallow.enabled = false`.

CDPD-66247: TestOzoneFileSystem.testListStatusOnKeyNameContainDelimiter is intermittent

None

OPSAPS-69539 : CDP Runtime 7.1.9 does not support Oracle JDK 8u401 or OpenJDK 1.8.0_402 (8u402). Some services will fail to start. This can be a problem on RHEL 9.x as version 8u402 is the default OpenJDK 8 installed by the OS.

Workaround is to install an earlier version of JDK 8. For example Oracle jdk-8u291 / 1.8.0_291, or OpenJDK 8u292 / 1.8.0_292.

A fresh install of 7.1.9 or its CHF's does not allow user to bypass the Setup Database screen for YARN Queue Manager

YARN Queue Manager in Cloudera Data Platform (CDP) Private Cloud Base 7.1.9 CHF 2 does not require you to install a PostGres database, therefore users should not see the Setup Database screen and should be able to skip the Setup Database screen. With this known issue, users who are conducting a fresh install of 7.1.9 or its CHF's are not able to bypass the Setup Database screen as expected.

1. When conducting a fresh install of YARN Queue Manager in 7.1.9 or its CHF's, you must ensure that you have both CDP and Cloudera Manager upgraded to 7.1.9 CHF 2.
2. When you reach the Setup Database screen in the Cloudera Manager installation wizard for Queue Manager, enter any dummy values for the following fields:
 - a. Database name: configstore
 - b. Database Username: dbuser
 - c. Database Password: dbpassword
 - d. Database Hostname: localhost

YARN Queue Manager will not connect to PostGres with the above details and will fall back to the embedded database.

3. Run the following script command in a browser console to enable the Continue button:

```
document.querySelector('.btn.next').removeAttribute('disabled');
```

4. Click Continue and proceed with the YARN Queue Manager installation.
5. After installation is complete, SSH into the host that has Queue Manager installed, and run this command: `sed -i 's/migrationCompleted=true/migrationCompleted=false' /var/lib/ha-doop-yarn/migration.properties`



Note: Enable Queue Manager in the YARN configurations, and restart YARN.

6. Restart YARN Queue Manager.

CDPD-61524: Ozone Storage Container Manager fails to start on upgrading from CDP Private Cloud Base 7.1.6 to 7.1.9 CHF1. Also, if you have upgraded from CDP Private Cloud Base 7.1.6 to 7.1.7 or 7.1.8 and then to 7.1.9, the upgrade fails.

None. Cloudera recommends you to reach out to the Support before performing the upgrade to CDP Private Cloud Base 7.1.9.

CDPD-62254: Ozone is not supported on SLES15 with CHF1.

If your cluster has Ozone, Cloudera recommends you to not upgrade to 7.1.9 CHF1.

CDPD-62834: Status of the deleted table is seen as ACTIVE in Atlas after the completion of navigator2atlas migration process

The status of the deleted table displays as ACTIVE.

None

CDPD-62837: During the navigator2atlas process, the hive_storagedesc is incomplete in Atlas

For the hive_storagedesc entity, some of the attributes are not getting populated.

None

CDPD-67651: Recon UI changes for Solr Health Check

When Solr is slow/down, Solr time to respond to Recon Heatmap query or sometimes does not respond at all. And the Recon heatmap does not load the heatmap data.

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet `atlas.authentication.method.file` is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

CDPD-79160: NPE while deleting BusinessMetadata

If business metadata is created without adding any applicable types, a `NullPointerException` is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

CDPD-66938: [Analyze] [Atlas] test_time_range tests fail

When the Apache Atlas server is running on a node which has time zone other than UTC, there might be a time of day when the search results might differ if the relative `CreateTime` date range filters of `TODAY`, `YESTERDAY`, etc. are used. 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.

Use explicit date range filters instead of using relative date range filters, such as, `TODAY`, `YESTERDAY`.

CDPD-76789: Creating tag with name description throws java.lang.ClassCastException

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

UnsupportedClassVersionError

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.

CDPD-79099: Unable to import hive tables using import script in Public Cloud

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.

Technical Service Bulletins**TSB 2025-820: Potential Data Integrity Issues Found in Ozone**

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

Knowledge article

For the latest update on this issue see the corresponding Knowledge 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

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.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)

Cumulative hotfix 7

You can review the list of cumulative hotfixes that were shipped for CDP Private Cloud Base version 7.1.9 CHF7.

Cloudera Runtime 7.1.9.14 (Cumulative Hotfix 7) download URL:

Parcel Repository Location

```
https://[USERNAME]:[PASSWORD]@archive.cloudera.com/p/cdh7/7.1.9.14/parcels/
```

Fixed issues in 7.1.9 CHF 7

Know more about the cumulative hotfix 7 for 7.1.9. This cumulative hotfix was released on June 6, 2024.



Note: Contact Cloudera Support for questions related to hotfixes.

Following is the list of fixes that were shipped for CDP Private Cloud Base version 7.1.9-1.cdh7.1.9.p14.53489573.

CDPD-69100: ozone sh key get displays NullPointerException after stopping two data hosts for EC.

This issue is now resolved.

A change in Ozone caused the links for the datanodes to not route through Knox on the Ozone SCM UI.

The datanodes links are now fixed and the links now correctly redirect to the correct Knox URL.

CDPD-67771: Terasort with S3A on FSO fails with the DIRECTORY_NOT_FOUND error.

This issue is now resolved.

KT-7540: KeyTrustee-KeyHSM - Upgraded protobuf-java to 3.16.3/3.19.6/3.20.3/3.21.7 due to CVE-2022-3171

Upgraded the protobuf-java version to 3.16.3/3.19.6/3.20.3/3.21.7 for KeyHSM due to CVE-2022-3171.

KT-7536: KeyTrustee-KeyHSM - Upgraded Jetty to 9.4.54.v20240208 due to CVE-2024-22201

Upgraded the Jetty version to 9.4.54.v20240208 due to CVE-2024-22201.

KT-7530: KeyTrustee-KeyHSM - Upgraded Spring Framework to 6.1.6/6.0.19/5.3.34 due to CVE-2024-22243, CVE-2024-22259 and CVE-2024-22262

Upgraded the Spring Framework version to 6.1.6/6.0.19/5.3.34 for KeyHSM due to CVE-2024-22243, CVE-2024-22259 and CVE-2024-22262.

CDPD-69563: [AUTOSYNC] HTTP Server fails to start with wildcard principal

Ozone role startup failed with the NoClassDefFoundError error, when the following conditions were met:

- Enable Kerberos Authentication for HTTP web consoles was checked
- Kerberos principal for HTTP access is set to *

This issue is now resolved.

CDPD-69425: [AUTOSYNC] Close SstFileReaderIterator in RocksDBCcheckpointDiffer

FileReader and FileReaderIterator did not close after the use and the system did not respond due to many open files. This issue is now resolved, and both FileReader and FileReaderIterator are now closed after use.

CDPD-69355: Exclude older versions of transitive libraries of jackson from Ranger-Kafka-plugin

Upgraded the Jackson Databind version to 2.15.0 and excluded older version of libraries.

CDPD-69306: Hue fails when trying to install npm16 on SLES 15

Upgraded npm16 to npm20.

CDPD-68924: [AUTOSYNC] Do not fail read of EC block if the last chunk is empty

Few EC blocks in a cluster had an empty final chunk. These blocks failed to read data and made the data unavailable, even when data was still present on disk. This issue is now resolved and the data is read even if there are empty final chunks in a block.

CDPD-68841: Ranger [7.1.9 chf7] - Upgraded Netty to 4.1.108.Final due to CVE-2024-29025

Upgraded the Netty version to 4.1.108 due to CVE-2024-29025.

CDPD-68807: Ranger - Upgraded Spring Framework to 6.1.6/6.0.19/5.3.34 due to CVE-2024-22243, CVE-2024-22259 and CVE-2024-22262

Upgraded the Spring Framework version to 5.3.34 due to CVE-2024-22243, CVE-2024-22259 and CVE-2024-22262.

CDPD-68736, CDPD-68737: Ranger - Upgraded Opensearch to 1.3.15 due to CVE-2023-45807

Upgraded Opensearch to 1.3.15 due to CVE-2023-45807.

CDPD-68723, CDPD-68724: [AUTOSYNC] EC Reconstruction does not issue put block to data index if it is unused

An issue in Apache Ozone resulted in the loss of certain user data blocks after EC reconstruction. This was triggered by the failure of a disk or node in the cluster. This issue is now resolved.

CDPD-68282: SMM UI - Upgraded Node JS version to 20.12.1 due to multiple CVEs

Upgraded the Node JS version to 20.12.1 due to multiple CVEs.

CDPD-68197: [Upgrade][EC] Reconstruction failing with "java.io.IOException: None of the block data have checksum"

EC reconstruction failed with the java.io.IOException: None of the block data have checksum error, when EC blocks were created without EC checksum data. This issue is now resolved.

CDPD-68193: Exclude services/roles from being discovered

New gateway-site.xml properties are added to exclude certain services and roles from being discovered during Cloudera Manager service discovery. The property names are gateway.cloudera.manager.service.discovery.excluded.service.types and gateway.cloudera.manager.service.discovery.excluded.role.types.

CDPD-68084: [7.1.x] KnoxCLI command for generating descriptor for a role type from a list of hosts

A new command is added to KnoxCLI to generate a topology descriptor from a list of URLs. The command usage/parameters are described in <https://github.com/apache/knox/pull/835>.

CDPD-68065: [7.1.x] - Add configurable socket / read timeout parameter to discovery client

The following gateway-site.xml properties are added for better control over Cloudera Manager service discovery timeouts in Knox:

- gateway.cloudera.manager.service.discovery.connect.timeout.ms
- gateway.cloudera.manager.service.discovery.connect.read.ms
- gateway.cloudera.manager.service.discovery.connect.write.ms

CDPD-67848: Rocks tools native lib should not be inside jar

Releng shipped the same jar for all the Operating System whichever the platform built first. In this case, having the library packed inside the jar caused incompatibilities. This issue is now resolved and the library is now packed outside the jar so that the native library loads from a particular path

CDPD-67608: SMM - Upgraded Jetty to 9.4.54.v20240208 due to CVE-2024-22201

Upgraded the Jetty version to 9.4.54.v20240208 due to multiple CVEs.

CDPD-67600: Knox - Upgraded PostgreSQL to 42.5.5/42.6.1/42.7.2 due to CVE-2024-1597

Upgraded the PostgreSQL version to 42.5.5/42.6.1/42.7.2 due to CVE-2024-1597.

CDPD-67568: Intermittent HTTP 401 error codes in Oozie tests due to Shiro unable to login: null

There was an intermittent HTTP 401 error in Oozie tests as org.glassfish.main.libpam4j produced errors. This issue is now resolved and dependency on org.kushuke.libpam4j is now switched back.

CDPD-67181: Ozone OM fail to start with a crash dump

Ozone Manager failed to start after an install of base cluster. This issue is now resolved

CDPD-67124: SMM - Upgraded Hibernate-Validator to 6.2.5.Final due to CVE-2023-1932

Upgraded the Hibernate-Validator version to 6.2.5 due to CVE-2023-1932.

CDPD-67113: [7.1.9] Backport KAFKA-13988: Mirrormaker 2 auto.offset.reset=latest not working

Streams Replication Manager (SRM) did not respond to auto.offset.reset=latest config. This issue is now resolved.

CDPD-67098: Ranger - Upgrade Commons-Compress to 1.26.0 due to CVE-2024-25710 and CVE-2024-26308

Upgraded the Commons-Compress version to 1.26.0 due to CVE-2024-25710 and CVE-2024-26308.

CDPD-66169: Hadoop - Upgraded bcpkix-jdk15on to 1.70+ due to CVE-2019-17359

Upgraded the bcpkix-jdk15on library to version 1.70 to address CVE-2019-17359.

CDPD-66048: [AUTOSYNC] Remove readLock from KeyValueHandler.checkContainerIsHealthy

Removed KeyValueHandler.checkContainerIsHealthy to improve Ozone data read performance.

CDPD-65490: The Knox topology did not update within 600 seconds

There was an issue when Knox compared the generated XML topology with the currently deployed XML topology, in highly concurrent environments. This issue is now resolved and a fix is made to the code, that parses the new topology.

CDPD-62837: [Atlas] [navigator2atlas] hive_storagedesc is incomplete in Atlas

Changes are made to the Nav2Atlas module to set the relationType as hive_table_storagedesc of relationship attribute table for every entity of hive_storagedesc.

CDPD-77283: SRM can skip records when replication-records-lag endoffset fetching fails

SRM no longer skips records during data replication when end offset fetching fails.

There are no Common Vulnerabilities and Exposures (CVE) fixed in this CHF.

Known issues in 7.1.9 CHF 7

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.1.9 CHF 7.

After upgrading the cluster from CDP Private Cloud Base 7.1.8 to CDP Private Cloud Base 7.1.9 and 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, ClassNotFoundException for org.cloudera.log4j.redactor.RedactorAppender class was identified. The error message is harmless as the upgrade is successful. The error existed previously and does not affect the Ozone service and its operation.

None.

When using S3A committer fs.s3a.committer.name=directory with fs.s3a.committer.staging.conflict-mode=replace to write to FSO buckets, the client fails with the following error.

```
DIRECTORY_NOT_FOUND
org.apache.hadoop.ozone.om.exceptions.OMException:
Failed to find parent directory of xxxxxxxx at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentID(OMFileRequest
at
org.apache.hadoop.ozone.om.request.file.OMFileRequest.getParentId(OMFileRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.om.request.s3.multipart.S3MultipartUploadCompleteRequest
at
org.apache.hadoop.ozone.protocolPB.OzoneManagerRequestHandler.handleWriteRequest
at
org.apache.hadoop.ozone.om.ratis.OzoneManagerStateMachine.runCommand(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 MPU. The problem is that 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.

CDPD-67095: In Knox on the Ozone SCM UI, the datanode links are not rewritten and they are not routing through Knox.

In Knox on the Ozone SCM UI the datanode links are not rewritten and they are not routing through Knox.

None.

CDPD-66508: Shallow listing is enabled by default in 7.1.9. There is a bug in shallow listing that causes the below error when listing an empty directory in a LEGACY/OBS bucket:

Error when listing an empty directory in a LEGACY/OBS bucket: mkdir: getFileStatus on s3a://testbucket/data/test: com.amazonaws.services.s3.model.AmazonS3Exception: Server Error (Service: Amazon S3; Status Code: 500; Error Code: 500 Server Error; Request ID: null; S3 Extended Request ID: null; Proxy: null), S3 Extended Request ID: null:500 Server Error: Server Error (Service: Amazon S3; Status Code: 500; Error Code: 500 Server Error; Request ID: null; S3 Extended Request ID: null; Proxy: null)

In S3 gateway log: Caused by: java.lang.IndexOutOfBoundsException: Index 0 out of bounds for length 0 at java.base/jdk.internal.util.Preconditions.outOfBounds(Preconditions.java:64) at java.base/jdk.internal.util.Preconditions.outOfBoundsCheckIndex(Preconditions.java:70) at java.base/jdk.internal.util.Preconditions.checkIndex(Preconditions.java:248) at java.base/java.util.Objects.checkIndex(Objects.java:372) at java.base/java.util.ArrayList.remove(ArrayList.java:535) at org.apache.hadoop.ozone.client.OzoneBucket\$KeyIterator.getNextShallowListOfKeys(OzoneBucket.java:1234) at org.apache.hadoop.ozone.client.OzoneBucket\$KeyIterator.getNextListOfKeys(OzoneBucket.java:1136) at org.apache.hadoop.ozone.client.OzoneBucket\$KeyIterator.hasNext(OzoneBucket.java:1110) at org.apache.hadoop.ozone.s3.endpoint.BucketEndpoint.get(BucketEndpoint.java:208) at jdk.internal.reflect.GeneratedMethodAccessor90.invoke(Unknown Source)

Disable shallow listing by performing the following steps:

1. Log in to Cloudera Manager
2. Navigate to Clusters
3. Select the Ozone service
4. Go to Configurations
5. In S3 Gateway Advanced Configuration Snippet (Safety Valve) for ozone-conf/ozone-site.xml, set ozone.s3g.list-keys.shallow.enabled = false.

CDPD-66247: TestOzoneFileSystem.testListStatusOnKeyNameContainDelimiter is intermittent

None

OPSAPS-69539 : CDP Runtime 7.1.9 does not support Oracle JDK 8u401 or OpenJDK 1.8.0_402 (8u402). Some services will fail to start. This can be a problem on RHEL 9.x as version 8u402 is the default OpenJDK 8 installed by the OS.

Workaround is to install an earlier version of JDK 8. For example Oracle jdk-8u291 / 1.8.0_291, or OpenJDK 8u292 / 1.8.0_292.

A fresh install of 7.1.9 or its CHF's does not allow user to bypass the Setup Database screen for YARN Queue Manager

YARN Queue Manager in Cloudera Data Platform (CDP) Private Cloud Base 7.1.9 CHF 2 does not require you to install a Postgres database, therefore users should not see the Setup Database screen and should be able to skip the Setup Database screen. With this known issue, users who are conducting a fresh install of 7.1.9 or its CHF's are not able to bypass the Setup Database screen as expected.

1. When conducting a fresh install of YARN Queue Manager in 7.1.9 or its CHF's, you must ensure that you have both CDP and Cloudera Manager upgraded to 7.1.9 CHF 2.

2. When you reach the Setup Database screen in the Cloudera Manager installation wizard for Queue Manager, enter any dummy values for the following fields:
 - a. Database name: configstore
 - b. Database Username: dbuser
 - c. Database Password: dbpassword
 - d. Database Hostname: localhost

YARN Queue Manager will not connect to PostGres with the above details and will fall back to the embedded database.

3. Run the following script command in a browser console to enable the Continue button:

```
document.querySelector('.btn.next').removeAttribute('disabled');
```

4. Click Continue and proceed with the YARN Queue Manager installation.
5. After installation is complete, SSH into the host that has Queue Manager installed, and run this command: `sed -i 's/migrationCompleted=true/migrationCompleted=false/' /var/lib/ha/hoop-yarn/migration.properties`



Note: Enable Queue Manager in the YARN configurations, and restart YARN.

6. Restart YARN Queue Manager.

CDPD-61524: Ozone Storage Container Manager fails to start on upgrading from CDP Private Cloud Base 7.1.6 to 7.1.9 CHF1. Also, if you have upgraded from CDP Private Cloud Base 7.1.6 to 7.1.7 or 7.1.8 and then to 7.1.9, the upgrade fails.

None. Cloudera recommends you to reach out to the Support before performing the upgrade to CDP Private Cloud Base 7.1.9.

CDPD-62254: Ozone is not supported on SLES15 with CHF1.

If your cluster has Ozone, Cloudera recommends you to not upgrade to 7.1.9 CHF1.

CDPD-62834: Status of the deleted table is seen as ACTIVE in Atlas after the completion of navigator2atlas migration process

The status of the deleted table displays as ACTIVE.

None

CDPD-62837: During the navigator2atlas process, the hive_storagedesc is incomplete in Atlas

For the hive_storagedesc entity, some of the attributes are not getting populated.

None

CDPD-67651: Recon UI changes for Solr Health Check

When Solr is slow/down, Solr time to respond to Recon Heatmap query or sometimes does not respond at all. And the Recon heatmap does not load the heatmap data.

CDPD-76035: Resource lookup for Atlas service is failing

Once the Atlas configuration snippet `atlas.authentication.method.file` is enabled and a classification is created, these do not synchronize correctly to the **Type Category** resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the **Type Name**.

CDPD-79160: NPE while deleting BusinessMetadata

If business metadata is created without adding any applicable types, a `NullPointerException` is produced when we try to delete that business metadata.

None

Apache Jira: [ATLAS-4863](#)

CDPD-66938: [Analyze] [Atlas] test_time_range tests fail

When the Apache Atlas server is running on a node which has time zone other than UTC, there might be a time of day when the search results might differ if the relative CreateTime date range filters of TODAY, YESTERDAY, etc. are used. 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.

Use explicit date range filters instead of using relative date range filters, such as, TODAY, YESTERDAY.

CDPD-76789: Creating tag with name description throws java.lang.ClassCastException

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-77435: RAZ: Import-hive on Cloudera Data Engineering cluster fails

Running import-hive command on RAZ enabled Cloudera Data Engineering cluster fails with error "NoClassDefFoundError: com/sun/jersey/core/spi/factory/ResponseBuilderImpl".

CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeUpException: null" is seen

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

UnsupportedClassVersionError

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.

CDPD-79099: Unable to import hive tables using import script in Public Cloud

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.

Technical Service Bulletins**TSB 2025-820: Potential Data Integrity Issues Found in Ozone**

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

Knowledge article

For the latest update on this issue see the corresponding Knowledge 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

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.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2025-835: Dry run of incremental Ozone replication can cause failure to replicate some changes in Cloudera Replication Manager](#)