

Cloudera Runtime 7.2.2

Release Notes

Date published: 2020-10-07

Date modified:

CLOUdera

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

Legal Notice

© Cloudera Inc. 2025. All rights reserved.

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

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

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

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

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

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

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

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

Contents

Overview.....	5
Cloudera Runtime Component Versions.....	5
Using the Cloudera Runtime Maven Repository.....	6
Maven Artifacts for Cloudera Runtime 7.2.2.....	6
What's New In Cloudera Runtime 7.2.2.....	25
What's New in Apache Atlas.....	25
What's new in Data Analytics Studio.....	25
What's New in Apache HBase.....	25
What's New in Apache Hadoop HDFS.....	25
What's New in Apache Hive.....	25
What's New in Hue.....	25
What's New in Apache Impala.....	25
What's New in Apache Kafka.....	27
What's New in Apache Knox.....	27
What's New in Apache Kudu.....	27
What's New in Apache Oozie.....	28
What's New in Apache Phoenix.....	28
What's New in Schema Registry.....	28
What's New in Cloudera Search.....	28
What's New in Apache Spark.....	29
What's New in Sqoop.....	29
What's new in Streams Messaging Manager.....	29
What's New in Apache Hadoop YARN.....	30
What's New in Apache ZooKeeper.....	30
Fixed Issues In Cloudera Runtime 7.2.2.....	30
Fixed Issues in Atlas.....	30
Fixed issues in Data Analytics Studio.....	32
Fixed Issues in Apache Hadoop.....	32
Fixed Issues in HBase.....	32
Fixed Issues in HDFS.....	32
Fixed Issues in Apache Hive.....	33
Fixed Issues in Hue.....	33
Fixed Issues in Apache Impala.....	33
Fixed Issues in Apache Kafka.....	33
Fixed Issues in Apache Knox.....	33
Fixed Issues in Apache Kudu.....	34
Fixed Issues in Apache Oozie.....	34
Fixed Issues in Phoenix.....	35
Fixed Issues in Apache Ranger.....	35
Fixed Issues in Schema Registry.....	36
Fixed Issues in Cloudera Search.....	36

Fixed Issues in Spark.....	37
Fixed Issues in Apache Sqoop.....	37
Fixed Issues in Streams Messaging Manager.....	38
Fixed Issues in Apache YARN.....	39
Fixed Issues in Zeppelin.....	39
Fixed Issues in Apache ZooKeeper.....	39

Known Issues In Cloudera Runtime 7.2.2..... 39

Known Issues in Apache Atlas.....	39
Known Issues in Data Analytics Studio.....	42
Known Issues in Apache HBase.....	43
Known Issues in HDFS.....	45
Known Issues in Apache Hive.....	45
Known Issues in Hue.....	48
Known Issues in Apache Impala.....	51
Known Issues in Apache Kafka.....	55
Known Issues in Kerberos.....	58
Known Issues in Apache Knox.....	58
Known Issues in Apache Kudu.....	58
Known Issues in Apache Oozie.....	58
Known Issues in Apache Phoenix.....	59
Known Issues in Apache Ranger.....	59
Known Issues in Schema Registry.....	59
Known Issues in Cloudera Search.....	59
Known Issues in Apache Solr.....	63
Known Issues in Apache Spark.....	63
Known Issues for Apache Sqoop.....	64
Known issues in Streams Messaging Manager.....	64
Known Issues in MapReduce and YARN.....	65
Known Issues in Apache Zeppelin.....	68
Known Issues in Apache ZooKeeper.....	69

Behavioral Changes In Cloudera Runtime 7.2.2..... 69

Behavioral Changes in Apache Phoenix.....	69
Behavioral Changes in Ranger.....	69
Behavioral Changes in Apache Hadoop YARN.....	70

Deprecation Notices In Cloudera Runtime 7.2.2..... 70

Deprecation notices in Apache Kudu.....	71
Deprecation Notices for Apache Kafka.....	71
Deprecation Notices in Apache HBase.....	71

Overview

You can review the Release Notes of Cloudera Runtime 7.2.2 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.

Cloudera Runtime Component Versions

You must be familiar with the versions of all the components in the Cloudera Runtime 7.2.2 distribution to ensure 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
Apache Atlas	2.0.0
Apache Calcite	1.19.0
Apache Avro	1.8.2
Apache Hadoop (Includes YARN and HDFS)	3.1.1
Apache HBase	2.2.5
Apache Hive	3.1.3000
Apache Impala	3.4.0
Apache Kafka	2.5.0
Apache Knox	1.3.0
Apache Kudu	1.12.0
Apache Livy	0.6.0
Apache MapReduce	3.1.1
Apache Oozie	5.1.0
Apache ORC	1.5.1
Apache Parquet	1.10.99
Apache Phoenix	5.0.0
Apache Ranger	2.0.0
Apache Solr	8.4.1
Apache Spark	2.4.5
Apache Sqoop	1.4.7
Apache Tez	0.9.1
Apache Zeppelin	0.8.2
Apache ZooKeeper	3.5.5

Other Components

Component	Version
Data Analytics Studio	1.4.2

Component	Version
GCS Connector	1.9.10
HBase Indexer	1.5.0
Hue	4.5.0
Search	1.0.0
Schema Registry	0.9.1
Streams Messaging Manager	2.1.0

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0
Hive Meta Store (HMS)	1.0.0
Hive on Tez	1.0.0
Hive Warehouse Connector	1.0.0
Spark Atlas Connector	0.1.0
Spark Schema Registry	1.1.0

Using the Cloudera Runtime Maven Repository

Information about using Maven to build applications with Cloudera Runtime components.

If you want to build applications or tools for use with Cloudera Runtime components and you are using Maven or Ivy for dependency management, you can pull the Cloudera Runtime artifacts from the Cloudera Maven repository. The repository is available at 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>
```

Maven Artifacts for Cloudera Runtime 7.2.2

The following table lists the project name, groupId, artifactId, and version required to access each RUNTIME artifact.

Project	groupId	artifactId	version
Apache Atlas	org.apache.atlas	atlas-authorization	2.0.0.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.atlas	atlas-aws-s3-bridge	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-classification-updater	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-client-common	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-client-v1	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-client-v2	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-common	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-distro	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-docs	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-graphdb-api	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-graphdb-common	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-graphdb-janus	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-index-repair-tool	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-intg	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-janusgraph-hbase2	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-notification	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-plugin-classloader	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-repository	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-server-api	2.0.0.7.2.2.0-244
	org.apache.atlas	atlas-testtools	2.0.0.7.2.2.0-244
	org.apache.atlas	hbase-bridge	2.0.0.7.2.2.0-244
	org.apache.atlas	hbase-bridge-shim	2.0.0.7.2.2.0-244
	org.apache.atlas	hbase-testing-util	2.0.0.7.2.2.0-244
	org.apache.atlas	hdfs-model	2.0.0.7.2.2.0-244
	org.apache.atlas	hive-bridge	2.0.0.7.2.2.0-244
	org.apache.atlas	hive-bridge-shim	2.0.0.7.2.2.0-244
	org.apache.atlas	impala-bridge	2.0.0.7.2.2.0-244
	org.apache.atlas	impala-bridge-shim	2.0.0.7.2.2.0-244
	org.apache.atlas	impala-hook-api	2.0.0.7.2.2.0-244
	org.apache.atlas	kafka-bridge	2.0.0.7.2.2.0-244
	org.apache.atlas	navigator-to-atlas	2.0.0.7.2.2.0-244
	org.apache.atlas	sqoop-bridge	2.0.0.7.2.2.0-244
	org.apache.atlas	sqoop-bridge-shim	2.0.0.7.2.2.0-244
Apache Avro	org.apache.avro	avro	1.8.2.7.2.2.0-244
	org.apache.avro	avro-compiler	1.8.2.7.2.2.0-244
	org.apache.avro	avro-ipc	1.8.2.7.2.2.0-244
	org.apache.avro	avro-mapred	1.8.2.7.2.2.0-244
	org.apache.avro	avro-maven-plugin	1.8.2.7.2.2.0-244
	org.apache.avro	avro-protobuf	1.8.2.7.2.2.0-244
	org.apache.avro	avro-service-archetype	1.8.2.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.avro	avro-thrift	1.8.2.7.2.2.0-244
	org.apache.avro	avro-tools	1.8.2.7.2.2.0-244
	org.apache.avro	trevni-avro	1.8.2.7.2.2.0-244
	org.apache.avro	trevni-core	1.8.2.7.2.2.0-244
Apache Calcite	org.apache.calcite	calcite-babel	1.19.0.7.2.2.0-244
	org.apache.calcite	calcite-core	1.19.0.7.2.2.0-244
	org.apache.calcite	calcite-druid	1.19.0.7.2.2.0-244
	org.apache.calcite	calcite-linq4j	1.19.0.7.2.2.0-244
	org.apache.calcite	calcite-server	1.19.0.7.2.2.0-244
	org.apache.calcite	calcite-ubenchmark	1.19.0.7.2.2.0-244
	org.apache.calcite.avatica	avatica	1.16.0.7.2.2.0-244
	org.apache.calcite.avatica	avatica-core	1.16.0.7.2.2.0-244
	org.apache.calcite.avatica	avatica-metrics	1.16.0.7.2.2.0-244
	org.apache.calcite.avatica	avatica-metrics-dropwizardmetrics	1.16.0.7.2.2.0-244
	org.apache.calcite.avatica	avatica-noop-driver	1.16.0.7.2.2.0-244
	org.apache.calcite.avatica	avatica-server	1.16.0.7.2.2.0-244
	org.apache.calcite.avatica	avatica-standalone-server	1.16.0.7.2.2.0-244
	org.apache.calcite.avatica	avatica-tck	1.16.0.7.2.2.0-244
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0.7.2.2.0-244
	org.apache.crunch	crunch-contrib	0.11.0.7.2.2.0-244
	org.apache.crunch	crunch-core	0.11.0.7.2.2.0-244
	org.apache.crunch	crunch-examples	0.11.0.7.2.2.0-244
	org.apache.crunch	crunch-hbase	0.11.0.7.2.2.0-244
	org.apache.crunch	crunch-hive	0.11.0.7.2.2.0-244
	org.apache.crunch	crunch-scrunch	0.11.0.7.2.2.0-244
	org.apache.crunch	crunch-spark	0.11.0.7.2.2.0-244
	org.apache.crunch	crunch-test	0.11.0.7.2.2.0-244
Apache Druid	org.apache.druid	druid-aws-common	0.17.1.7.2.2.0-244
	org.apache.druid	druid-benchmarks	0.17.1.7.2.2.0-244
	org.apache.druid	druid-console	0.17.1.7.2.2.0-244
	org.apache.druid	druid-core	0.17.1.7.2.2.0-244
	org.apache.druid	druid-gcp-common	0.17.1.7.2.2.0-244
	org.apache.druid	druid-hll	0.17.1.7.2.2.0-244
	org.apache.druid	druid-indexing-hadoop	0.17.1.7.2.2.0-244
	org.apache.druid	druid-indexing-service	0.17.1.7.2.2.0-244
	org.apache.druid	druid-integration-tests	0.17.1.7.2.2.0-244
	org.apache.druid	druid-processing	0.17.1.7.2.2.0-244
	org.apache.druid	druid-server	0.17.1.7.2.2.0-244
	org.apache.druid	druid-services	0.17.1.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.druid	druid-sql	0.17.1.7.2.2.0-244
	org.apache.druid	extendedset	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hadoop-extensions	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-security	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-bloom-filter	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-data-sketches	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-c2-extensions	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-google-extensions	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-hdfs-storage	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-histogram	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-kafka-extraction-namespace	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-kafka-indexing-service	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-kerberos	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-kinesis-indexing-service	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-lookups-cached-global	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-lookups-cached-single	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-orc-extensions	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-parquet-extensions	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-protobuf-extensions	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-s3-extensions	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-stats	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-sql-metadata-storage	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-sql-metadata-storage	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-ssl-client-sslcontext	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-emitter	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-emitter	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-emitter	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-extensions	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-storage	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-extensions	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-bount	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-extensions	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-emitter	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-sketch	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-average-query	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-emitter	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-cache	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-sketch	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-extensions	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-hbase-metrics-min-max	0.17.1.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.druid.extensions	druid-column-bloom	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-column-committer	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-column-contrib	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-column-view-maintenance	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-column-view-selection	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-column-storage	0.17.1.7.2.2.0-244
	org.apache.druid.extensions	druid-column-contrib	0.17.1.7.2.2.0-244
GCS Connector	com.google.cloud.bigtable	gcs-connector	2.1.2.7.2.2.0-244
	com.google.cloud.bigtable	gcs-connector	2.1.2.7.2.2.0-244
	com.google.cloud.bigtable	gcs-connector	2.1.2.7.2.2.0-244
	com.google.cloud.bigtable	gcs-connector	2.1.2.7.2.2.0-244
	com.google.cloud.bigtable	gcs-connector	2.1.2.7.2.2.0-244
Apache Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-annotations	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-archives	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-assemblies	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-auth	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-aws	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-azure	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-build-tools	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-client	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-client-api	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-common	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-datajoin	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-distcp	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-extras	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-fs2img	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-gridmix	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-hdds-client	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-hdds-common	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-hdds-config	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-hdds-container-service	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-hdds-docs	0.6.0.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-hdds-hadoop-dependency-client	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-hdds-hadoop-dependency-server	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-hdds-hadoop-dependency-test	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-hdds-interface-admin	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-hdds-interface-client	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-hdds-interface-server	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-hdds-server-framework	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-hdds-server-scm	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-hdds-test-utils	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-hdds-tools	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-hdfs	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-kafka	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-kms	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-minicluster	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-minikdc	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-nfs	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-openstack	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-client	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-common	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-csi	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-datanode	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-dist	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-filesystem	0.6.0.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-ozone-filessystem-common	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-filessystem-hadoop2	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-filessystem-hadoop3	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-filessystem-shaded	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-insight	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-integration-test	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-interface-client	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-network-tests	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-ozone-manager	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-recon	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-reconcodegen	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-s3gateway	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-tools	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-ozone-upgrade	0.6.0.7.2.2.0-244
	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-rumen	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-sls	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-streaming	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.2.2.0-244
	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.2.2.0-244
	org.apache.hadoop	mini-chaos-tests	0.6.0.7.2.2.0-244
Apache HBase	org.apache.hbase	hbase-annotations	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-checkstyle	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-client	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-client-project	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-common	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-endpoint	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-examples	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-external-blockcache	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-hadoop-compat	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-hadoop2-compat	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-hbtop	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-http	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-it	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-mapreduce	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-metrics	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-metrics-api	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-procedure	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-protocol	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-protocol-shaded	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-replication	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-resource-bundle	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-rest	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-rsgroup	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-server	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-shaded-client	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-shaded-client-project	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-shaded-mapreduce	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-shaded-testing-util	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-shaded-testing-util-tester	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-shell	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-testing-util	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-thrift	2.2.5.7.2.2.0-244
	org.apache.hbase	hbase-zookeeper	2.2.5.7.2.2.0-244
	org.apache.hbase.connector.kafka	hbase-connector-kafka-model	1.0.0.7.2.2.0-244
	org.apache.hbase.connector.kafka	hbase-connector-kafka-proxy	1.0.0.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.hbase.contrib	hbase-spark	1.0.0.7.2.2.0-244
	org.apache.hbase.contrib	hbase-spark-it	1.0.0.7.2.2.0-244
	org.apache.hbase.contrib	hbase-spark-protocol	1.0.0.7.2.2.0-244
	org.apache.hbase.contrib	hbase-spark-protocol-shaded	1.0.0.7.2.2.0-244
	org.apache.hbase.filesystem	hbase-cross	1.0.0.7.2.2.0-244
	org.apache.hbase.thirdparty	hbase-shaded-gson	3.3.0.7.2.2.0-244
	org.apache.hbase.thirdparty	hbase-shaded-miscellaneous	3.3.0.7.2.2.0-244
	org.apache.hbase.thirdparty	hbase-shaded-netty	3.3.0.7.2.2.0-244
	org.apache.hbase.thirdparty	hbase-shaded-protobuf	3.3.0.7.2.2.0-244
Apache Hive	org.apache.hive	hive-accumulo-handler	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-beeline	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-classification	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-cli	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-common	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-contrib	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-druid-handler	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-exec	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-hbase-handler	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-hplsql	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-jdbc	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-jdbc-handler	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-kryo-registrator	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-kudu-handler	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-llap-client	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-llap-common	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-llap-ext-client	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-llap-server	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-llap-tez	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-metastore	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-pre-upgrade	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-serde	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-service	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-service-rpc	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-shims	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-spark-client	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-standalone-metastore	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-storage-api	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-streaming	3.1.3000.7.2.2.0-244
	org.apache.hive	hive-testutils	3.1.3000.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.hive	hive-vector-code-gen	3.1.3000.7.2.2.0-244
	org.apache.hive	kafka-handler	3.1.3000.7.2.2.0-244
	org.apache.hive.hcatalog	hive-hcatalog-core	3.1.3000.7.2.2.0-244
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	3.1.3000.7.2.2.0-244
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	3.1.3000.7.2.2.0-244
	org.apache.hive.hcatalog	hive-hcatalog-streaming	3.1.3000.7.2.2.0-244
	org.apache.hive.hcatalog	hive-webhcat	3.1.3000.7.2.2.0-244
	org.apache.hive.hcatalog	hive-webhcat-java-client	3.1.3000.7.2.2.0-244
	org.apache.hive.shims	hive-shims-0.20	3.1.3000.7.2.2.0-244
	org.apache.hive.shims	hive-shims-0.23	3.1.3000.7.2.2.0-244
	org.apache.hive.shims	hive-shims-common	3.1.3000.7.2.2.0-244
	org.apache.hive.shims	hive-shims-scheduler	3.1.3000.7.2.2.0-244
Apache Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.2.2.0-244
Apache Kafka	org.apache.kafka	connect	2.5.0.7.2.2.0-244
	org.apache.kafka	connect-api	2.5.0.7.2.2.0-244
	org.apache.kafka	connect-basic-auth-extension	2.5.0.7.2.2.0-244
	org.apache.kafka	connect-file	2.5.0.7.2.2.0-244
	org.apache.kafka	connect-json	2.5.0.7.2.2.0-244
	org.apache.kafka	connect-mirror	2.5.0.7.2.2.0-244
	org.apache.kafka	connect-mirror-client	2.5.0.7.2.2.0-244
	org.apache.kafka	connect-runtime	2.5.0.7.2.2.0-244
	org.apache.kafka	connect-transforms	2.5.0.7.2.2.0-244
	org.apache.kafka	generator	2.5.0.7.2.2.0-244
	org.apache.kafka	jmh-benchmarks	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-clients	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-cloudera-plugins	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-examples	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-log4j-appender	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams-examples	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams-scala_2.12	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams-scala_2.13	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams-test-utils	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams-upgrade-system-tests-10	2.5.0.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.kafka	kafka-streams-upgrade-system-tests-11	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams-upgrade-system-tests-20	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams-upgrade-system-tests-21	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams-upgrade-system-tests-22	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams-upgrade-system-tests-23	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-streams-upgrade-system-tests-24	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka-tools	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka_2.12	2.5.0.7.2.2.0-244
	org.apache.kafka	kafka_2.13	2.5.0.7.2.2.0-244
Apache Knox	org.apache.knox	gateway-adapter	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-admin-ui	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-applications	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-cloud-bindings	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-demo-ldap	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-discovery-ambari	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-discovery-cm	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-docker	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-i18n	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-ha	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-jersey	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-rewrite	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-release	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-server	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-server-launcher	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-admin	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-as	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-definitions	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-hbase	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-health	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-hive	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-idbroker	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-impala	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-jkg	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-knoxsso	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-knoxssout	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-livy	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-metadata	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-nifi	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-rm	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-storm	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-test	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-tgs	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-vault	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-service-webhdfs	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-shell	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-shell-launcher	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-shell-release	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-shell-samples	1.3.0.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.knox	gateway-spi	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-test	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-test-idbroker	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-test-release-utils	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-test-utils	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-topology-simple	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-util-common	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-util-configinjector	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-util-launcher	1.3.0.7.2.2.0-244
	org.apache.knox	gateway-util-urltemplate	1.3.0.7.2.2.0-244
	org.apache.knox	hadoop-examples	1.3.0.7.2.2.0-244
	org.apache.knox	knox-cli-launcher	1.3.0.7.2.2.0-244
	org.apache.knox	knox-homepage-ui	1.3.0.7.2.2.0-244
	org.apache.knox	webhdfs-kerb-test	1.3.0.7.2.2.0-244
	org.apache.knox	webhdfs-test	1.3.0.7.2.2.0-244
Apache Kudu	org.apache.kudu	kudu-backup-tools	1.12.0.7.2.2.0-244
	org.apache.kudu	kudu-backup2_2.11	1.12.0.7.2.2.0-244
	org.apache.kudu	kudu-client	1.12.0.7.2.2.0-244
	org.apache.kudu	kudu-client-tools	1.12.0.7.2.2.0-244
	org.apache.kudu	kudu-hive	1.12.0.7.2.2.0-244
	org.apache.kudu	kudu-mapreduce	1.12.0.7.2.2.0-244
	org.apache.kudu	kudu-spark2-tools_2.11	1.12.0.7.2.2.0-244
	org.apache.kudu	kudu-spark2_2.11	1.12.0.7.2.2.0-244
	org.apache.kudu	kudu-test-utils	1.12.0.7.2.2.0-244
Apache Livy	org.apache.livy	livy-api	0.6.0.7.2.2.0-244
	org.apache.livy	livy-client-common	0.6.0.7.2.2.0-244
	org.apache.livy	livy-client-http	0.6.0.7.2.2.0-244
	org.apache.livy	livy-core_2.11	0.6.0.7.2.2.0-244
	org.apache.livy	livy-examples	0.6.0.7.2.2.0-244
	org.apache.livy	livy-integration-test	0.6.0.7.2.2.0-244
	org.apache.livy	livy-repl_2.11	0.6.0.7.2.2.0-244
	org.apache.livy	livy-rsc	0.6.0.7.2.2.0-244
	org.apache.livy	livy-scala-api_2.11	0.6.0.7.2.2.0-244
	org.apache.livy	livy-server	0.6.0.7.2.2.0-244
	org.apache.livy	livy-test-lib	0.6.0.7.2.2.0-244
	org.apache.livy	livy-thriftserver	0.6.0.7.2.2.0-244
	org.apache.livy	livy-thriftserver-session	0.6.0.7.2.2.0-244
Apache Lucene	org.apache.lucene	lucene-analyzers-common	8.4.1.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.lucene	lucene-analyzers-icu	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-analyzers-kuromoji	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-analyzers-morfologik	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-analyzers-nori	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-analyzers-openslp	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-analyzers-phonetic	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-analyzers-smartcn	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-analyzers-stempel	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-backward-codecs	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-benchmark	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-classification	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-codecs	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-core	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-demo	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-expressions	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-facet	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-grouping	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-highlighter	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-join	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-memory	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-misc	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-monitor	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-queries	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-queryparser	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-replicator	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-sandbox	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-spatial	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-spatial-extras	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-spatial3d	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-suggest	8.4.1.7.2.2.0-244
	org.apache.lucene	lucene-test-framework	8.4.1.7.2.2.0-244
Apache Oozie	org.apache.oozie	oozie-client	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-core	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-distro	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-examples	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-server	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-shaded-raz	5.1.0.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-sharelib-git	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-tools	5.1.0.7.2.2.0-244
	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.2.2.0-244
	org.apache.oozie.test	oozie-mini	5.1.0.7.2.2.0-244
Apache ORC	org.apache.orc	orc-core	1.5.1.7.2.2.0-244
	org.apache.orc	orc-examples	1.5.1.7.2.2.0-244
	org.apache.orc	orc-mapreduce	1.5.1.7.2.2.0-244
	org.apache.orc	orc-shims	1.5.1.7.2.2.0-244
	org.apache.orc	orc-tools	1.5.1.7.2.2.0-244
Apache Parquet	org.apache.parquet	parquet-avro	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-cascading	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-cascading3	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-column	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-common	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-encoding	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-format-structures	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-generator	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-hadoop	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-jackson	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-pig	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-pig-bundle	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-protobuf	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-scala_2.10	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-thrift	1.10.99.7.2.2.0-244
	org.apache.parquet	parquet-tools	1.10.99.7.2.2.0-244
Apache Phoenix	org.apache.phoenix	phoenix-client	5.0.0.7.2.2.0-244
	org.apache.phoenix	phoenix-core	5.0.0.7.2.2.0-244
	org.apache.phoenix	phoenix-hive	5.0.0.7.2.2.0-244
	org.apache.phoenix	phoenix-load-balancer	5.0.0.7.2.2.0-244
	org.apache.phoenix	phoenix-pherf	5.0.0.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.phoenix	phoenix-queryserver	5.0.0.7.2.2.0-244
	org.apache.phoenix	phoenix-queryserver-client	5.0.0.7.2.2.0-244
	org.apache.phoenix	phoenix-server	5.0.0.7.2.2.0-244
	org.apache.phoenix	phoenix-spark	5.0.0.7.2.2.0-244
	org.apache.phoenix	phoenix-tracing-webapp	5.0.0.7.2.2.0-244
Apache Ranger	org.apache.ranger	conditions-enrichers	2.0.0.7.2.2.0-244
	org.apache.ranger	credentialbuilder	2.0.0.7.2.2.0-244
	org.apache.ranger	embeddedwebserver	2.0.0.7.2.2.0-244
	org.apache.ranger	jisql	2.0.0.7.2.2.0-244
	org.apache.ranger	ldapconfigcheck	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-adls-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-atlas-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-atlas-plugin-shim	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-distro	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-examples-distro	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-hbase-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-hbase-plugin-shim	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-hdfs-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-hdfs-plugin-shim	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-hive-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-hive-plugin-shim	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-intg	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-kafka-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-kafka-plugin-shim	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-kms	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-kms-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-kms-plugin-shim	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-knox-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-knox-plugin-shim	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-kudu-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-kylin-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-kylin-plugin-shim	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-nifi-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-nifi-registry-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-ozone-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-ozone-plugin-shim	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-plugin-classloader	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-plugins-audit	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-plugins-common	2.0.0.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.ranger	ranger-plugins-cred	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-plugins-installer	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-raz-adls	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-raz-hook-abfs	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-raz-hook-s3	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-raz-intg	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-raz-processor	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-raz-s3	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-raz-s3-lib	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-rms-common	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-rms-hive	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-rms-plugins-common	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-rms-webapp	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-sampleapp-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-schema-registry-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-solr-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-solr-plugin-shim	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-sqoop-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-sqoop-plugin-shim	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-storm-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-storm-plugin-shim	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-tagsync	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-tools	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-util	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-yarn-plugin	2.0.0.7.2.2.0-244
	org.apache.ranger	ranger-yarn-plugin-shim	2.0.0.7.2.2.0-244
	org.apache.ranger	sample-client	2.0.0.7.2.2.0-244
	org.apache.ranger	sampleapp	2.0.0.7.2.2.0-244
	org.apache.ranger	shaded-raz-hook-abfs	2.0.0.7.2.2.0-244
	org.apache.ranger	unixauthclient	2.0.0.7.2.2.0-244
	org.apache.ranger	unixauthservice	2.0.0.7.2.2.0-244
	org.apache.ranger	unixusersync	2.0.0.7.2.2.0-244
Apache Solr	org.apache.solr	solr-analysis-extras	8.4.1.7.2.2.0-244
	org.apache.solr	solr-analytics	8.4.1.7.2.2.0-244
	org.apache.solr	solr-cell	8.4.1.7.2.2.0-244
	org.apache.solr	solr-clustering	8.4.1.7.2.2.0-244
	org.apache.solr	solr-core	8.4.1.7.2.2.0-244
	org.apache.solr	solr-dataimporthandler	8.4.1.7.2.2.0-244
	org.apache.solr	solr-dataimporthandler-extras	8.4.1.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.solr	solr-jaegertracer-configurator	8.4.1.7.2.2.0-244
	org.apache.solr	solr-langid	8.4.1.7.2.2.0-244
	org.apache.solr	solr-ltr	8.4.1.7.2.2.0-244
	org.apache.solr	solr-prometheus-exporter	8.4.1.7.2.2.0-244
	org.apache.solr	solr-security-util	8.4.1.7.2.2.0-244
	org.apache.solr	solr-solrj	8.4.1.7.2.2.0-244
	org.apache.solr	solr-test-framework	8.4.1.7.2.2.0-244
	org.apache.solr	solr-velocity	8.4.1.7.2.2.0-244
Apache Spark	org.apache.spark	spark-avro_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-catalyst_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-core_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-graphx_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-hadoop-cloud_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-hive_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-kubernetes_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-kvstore_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-launcher_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-mllib-local_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-mllib_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-network-common_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-network-shuffle_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-network-yarn_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-repl_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-shaded-raz	2.4.5.7.2.2.0-244
	org.apache.spark	spark-sketch_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-sql_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-streaming_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-tags_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-unsafe_2.11	2.4.5.7.2.2.0-244
	org.apache.spark	spark-yarn_2.11	2.4.5.7.2.2.0-244
Apache Sqoop	org.apache.sqoop	sqoop	1.4.7.7.2.2.0-244
	org.apache.sqoop	sqoop-test	1.4.7.7.2.2.0-244
Apache Tez	org.apache.tez	hadoop-shim	0.9.1.7.2.2.0-244
	org.apache.tez	hadoop-shim-2.8	0.9.1.7.2.2.0-244
	org.apache.tez	tez-api	0.9.1.7.2.2.0-244

Project	groupId	artifactId	version
	org.apache.tez	tez-aux-services	0.9.1.7.2.2.0-244
	org.apache.tez	tez-common	0.9.1.7.2.2.0-244
	org.apache.tez	tez-dag	0.9.1.7.2.2.0-244
	org.apache.tez	tez-examples	0.9.1.7.2.2.0-244
	org.apache.tez	tez-ext-service-tests	0.9.1.7.2.2.0-244
	org.apache.tez	tez-history-parser	0.9.1.7.2.2.0-244
	org.apache.tez	tez-javadoc-tools	0.9.1.7.2.2.0-244
	org.apache.tez	tez-job-analyzer	0.9.1.7.2.2.0-244
	org.apache.tez	tez-mapreduce	0.9.1.7.2.2.0-244
	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.2.2.0-244
	org.apache.tez	tez-runtime-internals	0.9.1.7.2.2.0-244
	org.apache.tez	tez-runtime-library	0.9.1.7.2.2.0-244
	org.apache.tez	tez-tests	0.9.1.7.2.2.0-244
	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.2.2.0-244
	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.2.2.0-244
	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.2.2.0-244
	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.2.2.0-244
Apache Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.2.2.0-244
	org.apache.zeppelin	zeppelin-display	0.8.2.7.2.2.0-244
	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.2.2.0-244
	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.2.2.0-244
	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.2.2.0-244
	org.apache.zeppelin	zeppelin-livy	0.8.2.7.2.2.0-244
	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.2.2.0-244
	org.apache.zeppelin	zeppelin-server	0.8.2.7.2.2.0-244
	org.apache.zeppelin	zeppelin-shaded-raz	0.8.2.7.2.2.0-244
	org.apache.zeppelin	zeppelin-shell	0.8.2.7.2.2.0-244
	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.2.2.0-244
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.5.5.7.2.2.0-244
	org.apache.zookeeper	zookeeper-client-c	3.5.5.7.2.2.0-244
	org.apache.zookeeper	zookeeper-contrib-loggraph	3.5.5.7.2.2.0-244
	org.apache.zookeeper	zookeeper-contrib-rest	3.5.5.7.2.2.0-244
	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.5.5.7.2.2.0-244
	org.apache.zookeeper	zookeeper-docs	3.5.5.7.2.2.0-244
	org.apache.zookeeper	zookeeper-jute	3.5.5.7.2.2.0-244
	org.apache.zookeeper	zookeeper-recipes-election	3.5.5.7.2.2.0-244
	org.apache.zookeeper	zookeeper-recipes-lock	3.5.5.7.2.2.0-244
	org.apache.zookeeper	zookeeper-recipes-queue	3.5.5.7.2.2.0-244

What's New In Cloudera Runtime 7.2.2

You must be aware of the additional functionalities and improvements to features of components in Cloudera Runtime 7.2.2. Learn how the new features and improvements benefit you.

What's New in Apache Atlas

There are no new features for Atlas in Cloudera Runtime 7.2.2.

What's new in Data Analytics Studio

There are no new features for Data Analytics Studio in Cloudera Runtime 7.2.2.

What's New in Apache HBase

Learn about the new features of HBase in Cloudera Runtime 7.2.2.

Access data on HBase using the Apache REST server

You can use the Apache HBase REST server to interact with the Apache HBase. You can use this as an alternative if you do not want to use the Java API. Interactions happen using URLs and the REST API. This feature is new for CDP public cloud deployments. For more information, see [Use the HBase REST server](#).

What's New in Apache Hadoop HDFS

There are no new features for Hadoop HDFS in Cloudera Runtime 7.2.2.

What's New in Apache Hive

There are no new features for Hive in Cloudera Runtime 7.2.2.

What's New in Hue

There are no new features for Hue in Cloudera Runtime 7.2.2.

What's New in Apache Impala

Learn about the new features of Impala in Cloudera Runtime 7.2.2.

Expanded SQL support

In this release, we extended SQL support in Impala to increase the completeness of SQL queries.

INTERSECT/EXCEPT

Added support for the set operators EXCEPT and INTERSECT. The EXCEPT and INTERSECT operators allow you to combine the result sets of multiple queries and return distinct rows by comparing the results of two queries.

See [UNION/INTERSECT/EXCEPT](#) for more information.

GROUPING SETS/ROLLUP/CUBE

Added support for GROUP BY ROLLUP, CUBE and GROUPING SETS. The GROUP BY ROLLUP clause creates a group for each combination of column expressions. The CUBE clause creates groups for all possible combinations of columns. The GROUPING SETS just lets you list out the combinations of expressions that you want to GROUP BY.

See [GROUP BY clause](#) for more information.

GROUPING and GROUPING_ID functions

Added these functions to be used in conjunction with the GROUP BY ROLLUP, CUBE, and GROUPING SETS clauses. GROUPING (column) function indicates whether the column in a GROUP BY list is aggregated or not. GROUPING_ID function computes the level of grouping and returns an integer value which is unique for each combination of grouping values.

See [GROUPING and GROUPING_ID functions](#) for more information.

Improved support for scalar subqueries

Uncorrelated scalar subqueries can now be used in SELECT list and in HAVING clause.

See [Subqueries in Impala SELECT statements](#) for more information.

New and Improved Performance Optimizations

The following changes have been made to improve the performance.

Pushdown limit to analytic sort operator

To optimize the performance, in 7.2.2, the Impala optimizer applies an outer LIMIT from a top level Sort down to the Sort below associated with an Analytic operator. This pushdown is considered if the Analytic operator is evaluating a ranking function such as RANK, DENSE_RANK, ROW_NUMBER and other qualifying criteria are met. Applying the limit early improves the performance by reducing the amount of data processed by the inner query. When the pushdown is applicable, the Analytic Sort is converted to a TopN Sort. Further, this is split into a bottom TopN and an upper TopN separated by a hash partition exchange. This ensures that the limit is applied as early as possible even before hash partitioning.

Example:

```
select * from (select l_partkey, l_quantity,
rank() over (partition by l_partkey order by l_quantity desc) rk
from lineitem) dt
where rk <= 100
order by l_partkey, l_quantity, rk
limit 100;
```

In this example, the limit 100 from the outer ORDER BY is pushed down to the analytic sort that is done below the Analytic operator.

Improved performance for sort

Improved the Sort operator efficiency. This reduces the execution time for top-n queries and partitioned loads.

Added distinct aggregation to semi joins

Added a new query option ENABLE_DISTINCT_SEMI_JOIN_OPTIMIZATION to reduce the number of potential plan regressions. When this query option is enabled the query planner considers inserting a distinct aggregation on the inner side of the join based on whether that aggregation would reduce the number of rows by more than 75%.

See [Query options](#) for more information.

Improved default runtime filter settings

The following runtime filter settings for DWX Impala have been updated to these new values to provide better results on DWX benchmarks. These are set as DWX defaults and can also be set via the default_query_options config:

- `RUNTIME_FILTER_MIN_SIZE=8192`
- `RUNTIME_FILTER_MAX_SIZE=2097152`
- `MAX_NUM_RUNTIME_FILTERS=50`
- `RUNTIME_FILTER_WAIT_TIME_MS=10000`

READ support for FULL-ACID ORC Tables

Until this release, Impala in CDP supported INSERT-ONLY transactional tables allowing both READ and WRITE operations. The latest version of Impala in CDP 7.2.2 now also supports READ of FULL ACID ORC tables without modifying any configurations.

See [READ support for FULL-ACID ORC Tables](#) for more information.

What's New in Apache Kafka

Learn about the new features of Kafka in Cloudera Runtime 7.2.2.

Detailed GC metrics are exported for brokers

Kafka brokers previously only exported an aggregated run count for garbage collection. From now on, more detailed information is exported. An aggregated time metric called `kafka_jvm_gc_time` is added. In addition, run count and time metrics are added for a handful of specific collectors which are the following:

- G1 Young Generation
- G1 Old Generation
- ParNew
- Concurrent Mark Sweep (CMS)
- PS MarkSweep
- PS Scavenge

The metrics for a particular collector are only exported if that collector is in use. For example, if you are using the G1 garbage collector, only the G1 collector metrics will be exported. Metrics for other collectors will be empty. The collector metrics are all named with the following naming convention `kafka_jvm_gc_<COLLECTOR_NAME>_<RUNS/TIME>`, for example `kafka_jvm_gc_parnew_runs`.

For a comprehensive list of available Kafka Broker metrics, see [Kafka Broker Metrics](#) in the Cloudera Manager Reference.

The kafka-clients library is no longer dependant on Apache Shiro and libpam4j

The `LdapPlainServerCallbackHandler` and `PamPlainServerCallbackHandler` custom Cloudera plugins introduced undesired dependencies for users of the `kafka-client` library. These plugins are now moved to a new module in Kafka that is only used on the broker side. As a result, the `kafka-clients` library is no longer dependent on Apache Shiro and `libpam4j`.

What's New in Apache Knox

There are no new features for Knox in Cloudera Runtime 7.2.2.

What's New in Apache Kudu

Learn about the new features of Kudu in Cloudera Runtime 7.2.2.

Table ownership

Kudu supports table ownership and can use the ownership to enforce owner authorization policies using Ranger. Ranger supports ownership privilege by creating a default policy that allows {OWNER} of a resource to access it without creating additional policy manually. A new access type, “delegate admin”, is added which grants the user permission to change ownership and create a table with a different owner.

KuduScanner.GetKuduTable method

The KuduTable instance is available from the KuduScanner, allowing the user to use the KuduTable instance populated by the scan token instead of making a GetTableSchema call to the master. The complete KuduTable is passed to the KuduScanner in the constructor, so it can be re-used by the consumer of the scanner. This eliminates additional roundtrip from client to master and the spikes of GetTableSchema requests to master whenever an Impala query or a Spark job starts on a large cluster.

Optimizations and improvements

- Adding kudu.snapshotTimestampMicros as an optional property to kudu spark readOptions. It allows consistent scans when timestamp is set before the first DataFrame read.
- The overall RPC performance increased by allocating larger chunks of memory from protobuf::Arena for all incoming calls. The boost in request rate is noticeable for RPCs with larger request and responses (for example, up to 25% increase is shown in synthetic benchmarks for GetTableLocations), but not so much for RPCs with smaller requests and responses (for example, no measurable gain in request rate for GetTableSchema).
- A load meter is introduced for ThreadPool, aiming to use active queue management techniques (AQM) such as CoDel in scenarios where thread pool queue load metrics are applicable.

What's New in Apache Oozie

There are no new features for Oozie in Cloudera Runtime 7.2.2.

What's New in Apache Phoenix

There are no new features for Phoenix in Cloudera Runtime 7.2.2.

What's New in Schema Registry

Learn about the new features of Schema Registry in Cloudera Runtime 7.2.2.

Security-related header controls in Schema Registry responses

The following HTTP headers appear in Schema Registry HTTP responses to provide security-related restrictions and instructions to the browser:

- Cache-Control
- Content-Security-Policy
- Content-Type-Options
- Strict-Transport-Security (if SSL is enabled)
- X-Frame options
- XSS-Protection

What's New in Cloudera Search

Learn about the new features of Cloudera Search in Cloudera Runtime 7.2.2.

Spark-Solr connector

With the Spark-Solr connector you can perform ETL batch and streaming jobs on large data sets, indexing data into Solr in multiple ways. Using the SparkApp framework you get the robustness of Spark, while you can access Solr easily and in a familiar way. The framework supports using both Scala and Java.

To learn more about batch indexing data into Solr using the connector, see [Batch indexing to Solr using SparkApp framework in spark-submit](#).

What's New in Apache Spark

Learn about the new features of Spark in Cloudera Runtime 7.2.2.

Data engineering cluster

You can create a data engineering cluster in Amazon AWS from within CDP by selecting the Data Engineering cluster template. A data engineering includes Spark, Livy, Hive, Zeppelin, and Oozie, along with supporting services (HDFS, YARN, and Zookeeper).

See [Creating a Cluster on AWS](#).

What's New in Sqoop

Learn what's new in the Apache Sqoop client in Cloudera Runtime 7.2.2.

To access the latest Sqoop documentation on Cloudera's documentation web site, go to [Sqoop Documentation 1.4.7.7.1.6.0](#).

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.

What's new in Streams Messaging Manager

Learn about the new features of Streams Messaging Manager in Cloudera Runtime 7.2.2.

Security-related header controls in SMM REST API server responses

The following HTTP headers appear in SMM REST API HTTP responses to provide security-related restrictions and instructions to the browser:

- Strict-Transport-Security
- Cache-Control

Security-related header controls in SMM UI server responses

The following HTTP headers appear in SMM UI server HTTP responses to provide security-related restrictions and instructions to the browser:

- Cache-Control

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

What's New in Apache Hadoop YARN

Learn about the new features of Hadoop YARN in Cloudera Runtime 7.2.2.

Resource Allocation - Absolute Mode

You can now allocate resources to queues using Absolute Mode. You can specify the number of units of vCores and memory resources by using the Absolute allocation mode.

For more information about resource allocation mode, see [Resource Allocation Overview](#).

What's New in Apache ZooKeeper

There are no new features for ZooKeeper in Cloudera Runtime 7.2.2.

Fixed Issues In Cloudera Runtime 7.2.2

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.2.

Fixed Issues in Atlas

Review the list of Apache Atlas issues that are resolved in Cloudera Runtime 7.2.2.

CDPD-1138: Spark Atlas Connector tracks column-level lineage

This issue is now resolved.

CDPD-11790: Shell entity is not resolved to the Complete entity under certain conditions.

Shell entities with duplicate qualifiedName are no longer created when processing message from Spark Atlas Connector. This issue is now resolved.

CDPD-14031: In the Spark Atlas Connector, few S3 entities are created using the V1 S3 model instead of the updated V2 S3 model.

Use Atlas S3 v2 models in Spark Atlas Connector. This issue is now resolved.

OPSAPS-57947: Kafka Broker SSL configuration is not correct in High Availability mode.

When deploying the DataHub in High Availability mode, some of the Ranger and Atlas configurations are not computed correctly. In particular atlas.kafka.security.protocol in Atlas, the SSL properties and the REST URL services depending on Ranger.

CDPD-13645: Contains sortBy=name, 'name' attribute is not in hive_storagedesc definition. Also if sortBy is not passed, default attribute is name

Validated if sortBy attribute passed in the request is present in relationship end definition, if not present, ignore sorting.

Validated if sortBy attribute is not passed, default attribute name is present in relationship end definition, if not present, ignore sorting.

CDPD-10873

1) Fixed quick search aggregation metrics when filtered with System Attributes

- 2) Fixed quick search aggregation metrics when filtering with more than one filter
- 3) Fixed quick search aggregation metrics when filtering with negation operator

CDPD-13805: Relationship api request will have provision to specify attributes to be present in search result.

Example Request: /v2/search/relationship?guid=ac9e04cc-f927-4334-af08-c83bc3733f5b&relation=columns&sortBy=name&sortOrder=ASCENDING&attributes=dcProfiledData

CDPD-11681:

1. Filter Search Results with multiple entity type by 'comma' separated string of typeName in the request Eg. "typeName": "hive_table,hive_db".
2. Filter Search Results with multiple tag by 'comma' separated string of tags in the request Eg. "classification": "tag1,tag2".

CDPD-13199: Incorrect attribute values in bulk import

When importing Business Metadata attribute assignments, Atlas used only the last assigned attribute value instead of individual values for each entity in the import list.

CDPD-372: All Spark Queries from the Same Spark Session were included in a Single Atlas Process

A Spark session can include multiple queries. When Atlas reports the Spark metadata, it creates a single process entity to correspond to the Spark session. The result was that an Atlas lineage picture showed multiple input entities or multiple output entities for a process, but the inputs and outputs were only related by the fact that they were included in operations in the same Spark session. In this release, the Spark Atlas Connector produces a spark_application entity for each Spark job. Each data flow produced by the job creates a spark_process entity in Atlas, which tracks the actual input and output data sets for that process. For more information, see [Spark metadata collection](#).

CDPD-12620: Migration progress bar not refreshed

During the import stage of Cloudera Navigator to Apache Atlas migration, the migration progress bar does not correctly refresh the migration status. The Statistics page in the Atlas UI displays the correct details of the migration.

This issue is resolved.

CDPD-10151: Short Spark job processes may be lost

In rare occasions, it is possible for events captured for Atlas by the Spark Atlas Connector to be dropped before the metadata reaches Atlas. It is more likely that an event is lost in very short running jobs.

This issue is resolved.

CDPD-6042: Hive Default Database Location Incorrect in Atlas Metadata

The location of the default Hive database as reported through the HMS-Atlas plugin does not match the actual location of the database. This problem does not affect non-default databases.

This issue is resolved.

CDPD-4662: Lineage graph links not working

Atlas lineage graphs do not include hyperlinks from assets to the assets' detail pages and clicking an asset does not provide an error in the log. Clicking an edge in a graph still provides access to edge behavior options such as controlling how classifications propagate.

This issue is resolved.

CDPD-3700: Missing Impala and Spark lineage between tables and their data files

Atlas does not create lineage between Hive tables and their backing HDFS files for CTAS processes run in Impala or Spark.

This issue is resolved.

Additional Cloudera JIRAs: CDP-5027, CDPD-3700, and IMPALA-9070

Fixed issues in Data Analytics Studio

There are no fixed issues for Data Analytics Studio in Cloudera Runtime 7.2.2.

Fixed Issues in Apache Hadoop

Review the list of Hadoop issues that are resolved in Cloudera Runtime 7.2.2.

CDPD-7383: The ABFS and ADL connectors compatible with Alpine Linux and other platforms which have libssl1.1-1.1.1b-r1 as their native OpenSSL implementation

See [HADOOP-16460](#) and [HADOOP-16438](#)

CDPD-15133: HADOOP-17130: Configuration.getValByRegex() should not update the results while fetching as it can cause ConcurrentModificationException.

This issue is now resolved.

Fixed Issues in HBase

Review the list of HBase issues that are resolved in Cloudera Runtime 7.2.2.

OPSAPS-57394: Create new Cloudera Manager metrics for HBase 2.0 JMX RIT metrics

The following HBase metrics are available in Cloudera Manager now:

- regions_in_transition_duration_num_ops
- regions_in_transition_duration_min
- regions_in_transition_duration_max
- regions_in_transition_duration_mean
- regions_in_transition_duration_25th_percentile
- regions_in_transition_duration_median
- regions_in_transition_duration_75th_percentile
- regions_in_transition_duration_90th_percentile
- regions_in_transition_duration_95th_percentile
- regions_in_transition_duration_98th_percentile
- regions_in_transition_duration_99th_percentile
- regions_in_transition_duration_99_9th_percentile

Fixed Issues in HDFS

Review the list of HDFS issues that are resolved in Cloudera Runtime 7.2.2.

CDPD-6100

This improvement makes HDFS NameNode leave safemode immediately if all blocks have reported in, reducing the cluster startup time.

CDPD-2946: Slow reading and writing of erasure-coded files

The ISA-L library is not packaged with HDFS as a result of which HDFS erasure coding falls back to the Java implementation which is much slower than the native Hadoop implementation. This slows down the reading and writing of erasure-coded files.

OPSAPS-43909: Execution filter is not applied to Delete Policy.

Execution filter is now applied to Delete Policy also. This issue is now resolved.

Fixed Issues in Apache Hive

Review the list of Hive issues that are resolved in Cloudera Runtime 7.2.2.

CDPD-14820: Fixed probeDecode issue. TezCompiler pushes down MapJoin Operators with Key expressions.

Now, only MapJoins with simple keys is supported by the Hive probeDecode feature.

OPSAPS-57720: DH cluster creation fails post CDH upgrade.

Cloudera Manager had no upgrade handler for Hive Metastore. This issue is now resolved.

OPSAPS-49148: "Update Hive Metastore NameNodes" invokes metatool for each database.

Removed unnecessary executions of the metatool with updateLocation to lower the total execution time and usage of resources for 'Update Hive Metastore NameNodes'. This issue is now resolved.

OPSAPS-43909: Execution filter is not applied to Delete Policy.

Execution filter is now applied to Delete Policy also. This issue is now resolved.

Fixed Issues in Hue

There are no fixed issues for Hue in Cloudera Runtime 7.2.2.

Fixed Issues in Apache Impala

Review the list of Impala issues that are resolved in Cloudera Runtime 7.2.2.

CDPD-14548: Upgrades a library - PostgreSQL JDBC Driver (pgjdbc) bundled with Impala. The previous version 42.2.5 was affected with vulnerability CVE-2020-13692. It is upgraded to version 42.2.14.

This issue is now resolved.

CDPD-10444: Update the version of Atlas used by Impala

Resolves the inconsistency between the versions of jackson-databind used by Atlas and Impala so that you no longer see the exception NoClassDefFoundError after the call to QueryEventHookManager#executeQueryCompleteHooks() in Impala.

Fixed Issues in Apache Kafka

Review the list of Kafka issues that are resolved in Cloudera Runtime 7.2.2.

KAFKA-9839: IllegalStateException on metadata update when broker learns about its new epoch after the controller

This is a backported fix. The broker now accepts control requests if the broker epoch is newer than the current epoch.

Fixed Issues in Apache Knox

Review the list of Knox issues that are resolved in Cloudera Runtime 7.2.2.

OPSAPS-57448: IDBroker does not export correct RDC configuration in HA.

The RDC configs is now correctly exported when IDBroker is in HA mode. This issue is now resolved.

Technical Service Bulletins

TSB 2022-553: DOM based XSS Vulnerability in Apache Knox

For the latest update on this issue see the corresponding Knowledge article: [TSB 2022-553: DOM based XSS Vulnerability in Apache Knox \(“Knox”\)](#)

Fixed Issues in Apache Kudu

Review the list of Kudu issues that are resolved in Cloudera Runtime 7.2.2.

KUDU-2373: maintenance-manager-num-threads=0 causes CHECK fail during server startup

Setting the gflag maintenance_manager_num_threads to 0 causes a failure during server setup. A flag validator was added to prevent this.

KUDU-3184: Fix GetFQDN() when canonical name returns null

If FQDN did not have domain name (ex .local), canonical name was not assigned and FQDN returned null.

Fixed Issues in Apache Oozie

Review the list of Oozie issues that are resolved in Cloudera Runtime 7.2.2.

CDPD-14964: When a Java action called System.exit() that resulted in a misleading security exception for Sqoop actions, there was an error because of the misleading exception in the yarn logs even though the Workflow is successful.

This issue is now resolved.

CDPD-15735: Oozie Spark actions are failing because Spark and Kafka are using different Scala versions.

This issue is now resolved.

OPSAPS-57429: Zookeeper SSL/TLS support for Oozie.

When SSL is enabled in Zookeeper, Oozie tries to connect to Zookeeper using SSL instead of a non-secure connection.

CDPD-14600

Apache ActiveMQ is updated to address CVE-2016-3088

CDPD-13702

The PostgreSQL driver is upgraded to address CVE-2020-13692

CDPD-11967

Fix to address CWE-693: Protection Mechanism Failure

CDPD-12742: Oozie was not able to communicate with ID Broker and hence it failed to obtain a delegation token, because of a missing Jar

That Jar is now deployed together with Oozie and hence the underlying issue is fixed.

CDPD-12283: By Oozie did not allow to use s3a and abfs file systems and users had to manually specify the supportability of these via Safety Valve

Since Oozie is compatible with these filesystems we changed the default Oozie configuration to allow these so users don't have to manually specify it.

CDPD-10746: Fix to address CVE-2019-17571

CDPD-9895: Various errors when trying to use an S3 filesystem

Oozie is now fully compatible with S3.

CDPD-9761: There is a sub workflow run in independent mode that runs a fork action which contains two (or more) actions

These actions inside the fork action run in parallel mode, and they have some seconds delay in between them. If a parameter is passed to one of these actions, that cannot be resolved, then it changes its status to FAILED, and also the workflows state to FAILED. The other actions state

which are not started yet will stuck in PREP state forever. The correct behaviour would be to KILL the remaining actions as well as the workflow. Note: this bug only occurs when it is run in independent mode. If it has a parent workflow, then the parent workflow will kill this workflow after 10 minutes because of the callback process.

CDPD-9721: Upgrade built-in spark-hive in Oozie

Oozie is using the Spark-Hive library from the stack.

CDPD-9220: Oozie spark actions using --keytab fail due to duplicate dist. cache

Oozie spark actions add everything in the distributed cache of the launcher job to the distributed cache of the spark job, meaning the keytab is already there, then the --keytab argument tries to add it again causing the failure.

CDPD-9189: Apache Pig support was completely removed from Oozie

CDPD-7108: In case we have a workflow which has, lets say, 80 actions after each other, then the validator code "never" finishes

CDPD-7107: The following were added to the spark opts section of the spark action: --conf spark

CDPD-7106: query tag is not functional for Hive2 action node in oozie

Workflow is intended to create a hive table using Hive2 action node. Though workflow run successfully, table is not created.

CDPD-7105: Oozie workflow processing becomes slow after the increase of rows in WF_JOBS and WF_ACTIONS tables when running against SQL Server

CDPD-6877: When you create a MapReduce action which then creates more than 120 counters, an exception was thrown

CDPD-6630: Oozie by default gathers delegation tokens for the nodes defined in MapReduce

CDPD-5168: Logging enhancements in CoordElFunctions for better supportability

CDPD-4826: Oozies web server does not work when TLS is enabled and Open JDK 11 is in use

This issue is now fixed.

Fixed Issues in Phoenix

There are no fixed issues for Phoenix in Cloudera Runtime 7.2.2.

Fixed Issues in Apache Ranger

Review the list of Ranger issues that are resolved in Cloudera Runtime 7.2.2.

CDPD-15401: When you enable Hive Metastore lookup in Ranger admin, resource lookup returns nothing and displays an error.

This issue is now resolved. You must use this step as a solution: sudo

```
ln -s /opt/cloudera/parcels/*<CDH-version>*/jars/libfb303-0.9.3.jar  
/opt/cloudera/parcels/*<CDH-version>*/lib/ranger-admin/ews/webapp/WEB-INF/lib/libfb303-0.9.3.jar
```

CDPD-14269 and CDPD-14289: Failed resource filtering in Ranger Policy Export.

Exporting tag policies result in a 204 error when the polResource query parameter is used.

CDPD-12848: When you try to create multiple policies using the API having same non-existing group, the group creation fails due to multiple threads trying to create the same group at once.

Separate threads are now created for retry group creation and checks if the group is previously created and associate it with policy.

CDPD-10072: Ranger Ozone plugin unable to write to solr audits in SSL enabled cluster

This issue is now resolved. A separate folder libext is added under the Ozone library path and all the ranger plugin jars are added under this new folder.

OPSAPS-57495: The Ranger role-level principal for Ranger Admin, Ranger Usersync, and Ranger Tagsync can now be customized from the Cloudera Manager UI.

Fixed Issues in Schema Registry

Review the list of Schema Registry issues that are resolved in Cloudera Runtime 7.2.2.

CDPD-15097: Schemas with null description do not show up in the UI.

Previously, when using the UI or API calls, search results would sometimes list schemas according to description instead of name matches. Schemas without a description would not appear in the results.

In this release, schemas with a null description now appear and the UI search results are listed correctly.

CDPD-15066: Implement schema version selection in Ranger.

Users could not apply a Ranger policy to a specific schema version. This was because the schema version field on the *Create* or *Edit Policy* page did not accept a specific version and instead returned an asterisk. The asterisk indicated that the policy would apply to all versions of the schema.

Users can now apply Ranger policies to specific schema versions.

OPSAPS-57467: Hard-coded parcel directory causes failure when non-standard path is used and TLS is enabled.

The parcel directory path was hard-coded in the Schema Registry scripts and caused failures when a non-standard path was used and TLS was enabled.

The parcel directory path is no longer hard-coded in the Schema Registry CSD.

OPSAPS-57409: Add security-related header controls to all Schema Registry responses

The following HTTP headers appear in Schema Registry HTTP responses to provide security-related restrictions and instructions to the browser:

- Cache-Control
- Content-Security-Policy
- Content-Type-Options
- Strict-Transport-Security (if SSL is enabled)
- X-Frame options
- XSS-Protection

OPSAPS-57294: Schema Registry first run fails when multiple Ranger Admin services are configured.

The first run of Schema Registry failed when more than one Ranger Administrator service was configured.

The Schema Registry start-up script can now handle multiple Ranger Administrator services.

Fixed Issues in Cloudera Search

Review the list of Cloudera Search issues that are resolved in Cloudera Runtime 7.2.2.

CDPD-16332

Solr - Backport fixes for CVE-2020-13941 - SOLR-14515 / SOLR-14561

Fixed Issues in Spark

Review the list of Spark issues that are resolved in Cloudera Runtime 7.2.2.

CDPD-1138: Spark Atlas Connector tracks column-level lineage

This issue is now resolved.

CDPD-14906: Spark reads or writes TIMESTAMP data for values before the start of the Gregorian calendar. This happens when Spark is:

- Using dynamic partition inserts.
- Reading or writing from an ORC table when `spark.sql.hive.convertMetastoreOrc=false` (the default is true).
- Reading or writing from an Orc table when `spark.sql.hive.convertMetastoreOrc=true` but `spark.sql.orc.impl=hive` (the default is native).
- Reading or writing from a Parquet table when `spark.sql.hive.convertMetastoreParquet=false` (the default is true).

This issue is now resolved.

CDPD-15385: Currently, delegation token support for Spark DStreams is not available.

Added Kafka delegation token support for DStreams in the Spark 2.4.5. This issue is now resolved.

CDPD-15735: Oozie Spark actions are failing because Spark and Kafka are using different Scala versions.

This issue is now resolved.

CDPD-10532: Update log4j to address CVE-2019-17571

Replaced log4j with an internal version to fix CVE-2019-17571.

CDPD-10515: Incorrect version of jackson-mapper-asl

Use an internal version of jackson-mapper-asl to address CVE-2017-7525.

CDPD-7882: If an insert statement specifies partitions both statically and dynamically, there is a potential for data loss

To prevent data loss, this fix throws an exception if partitions are specified both statically and dynamically. You can follow the workarounds provided in the error message.

CDPD-15773: In the previous versions, applications that share a Spark Session across multiple threads was experiencing a deadlock accessing the HMS.

This issue is now resolved.

Technical Service Bulletins

TSB 2021-441: CDP Powered by Apache Spark may incorrectly read/write pre-Gregorian timestamps

For the latest update on this issue see the corresponding Knowledge article: [TSB-2021 441: Spark may incorrectly read/write pre-Gregorian timestamps](#)

Fixed Issues in Apache Sqoop

Review the list of Sqoop issues that are resolved in Cloudera Runtime 7.2.2.

CDPD-12646: Sqoop does not close the open database connection before submitting the MapReduce Job. The open connection utilizes resources and displays an error message in the log when the connection times out.

This issue is resolved.

Fixed Issues in Streams Messaging Manager

Review the list of Streams Messaging Manager issues that are resolved in Cloudera Runtime 7.2.2.

CDPD-12981: Security related header control is not available in Streams Messaging Manager.

Added Security related HTTP headers to Streams Messaging Manager UI server responses: Cache Control, XSS-PROTECTION, Content-Security-Policy, Strict-Transport-Security, X-Content-Type-Options, and X-Frame-Options.

OPSAPS-57410: Added Security related headers to Streams Messaging Manager Rest API Server responses.

Strict-Transport-Security and Cache-Control are added.

CDPD-12147: SMM throws NumberFormatException error.

SMM throws a NumberFormatException error when the SMM's JVM is set to non-US Locale.

CDPD-14383: SMM calls CM time series API less times.

SMM calls Cloudera Manager time series API less times because of grouping metric requests into one call. This can be turned off by setting the following property in kafkaMetricsConfig properties:

request.metrics.separately: true

CDPD-14544: SMM displays smaller topic metric values for larger time.

On the overview page, smaller topic metric values were displayed for larger time periods in comparison to the smaller time periods due to larger rollup intervals.

CDPD-14930: SMM displays incorrect values of metrics.

For certain time periods, metrics such as Bytes In, Bytes out, Messages In are displaying incorrect value. This is due to CM metrics rollup where granularity of more than 1 min causes metrics to lose metric information resulting in incorrect calculation.

CDPD-15552: SMM throws ConsumerGroups exception.

SMM tries to parse Kafka connectors as ConsumerGroups and throws the following exception:

```
Error reading field 'X': Error reading string of length....
```

CDPD-15738: Knox directs to wrong SMM address.

Knox is directing to the wrong address for SMM.

OPSAPS-57411: SMM metrics fetching mode not configurable.

SMM metrics fetching mode is not configurable. By default it is calling CM API timeseries endpoint with multiple metrics separated by comma.

New configuration parameter is introduced in SMM to make it possible to control metrics fetching mode.

OPSAPS-57444: Keystore and private key password mismatch.

SMM throws an error if the keystore and private key password are not the same.

OPSAPS-57468: Hardcoded parcel directory causes failure.

The parcel directory path was hard-coded in the SMM scripts and caused failures when a non-standard path was used and TLS was enabled.

The parcel directory path is no longer hard-coded.

OPSAPS-57539: SMM not usable after upgrade.

When upgrading the Cloudera Runtime, the previous Runtime's SMM UI process was not being killed and making SMM unusable.

Now, the SMM UI stop script has been improved, so that it will kill the child processes to prevent the SMM UI process from being stuck.

OPSAPS-57745: CM displays incorrect role.

Cloudera Manager displays incorrect role status when the SMM UI process fails or stops.

Fixed Issues in Apache YARN

Review the list of YARN issues that are resolved in Cloudera Runtime 7.2.2.

OPSAPS-56456: Application history is lost for Mapreduce applications after the upgrade.

The log aggregation file controllers suffix configurations are automatically changed during the upgrade to a CDP cluster. This issue is now resolved.

OPSAPS-27702: YARN configuration for Limit Nonsecure Container Executor Users does not function as expected.

Created a new upgradehandler which copies the SV value to the newly introduced ParamSpec during the upgrade. This issue is now resolved.

Fixed Issues in Zeppelin

Review the list of Zeppelin issues that are resolved in Cloudera Runtime 7.2.2.

CDPD-10187: Incorrect version of jackson-mapper-asl.

Use internal version of jackson-mapper-asl to handle CVE-2017-7525.

CDPD-1683: Zeppelin demo users have been removed

Use cluster users to access Zeppelin. For information on provisioning users in CDP, see [Onboarding users](#).

CDPD-880, CDPD-1685: Shell, JDBC, and Spark interpreters have been removed

Use an available interpreter. For Spark functionality, use the Livy interpreter.

CDPD-3047: Markdown interpreter does not handle certain numbered list syntax correctly

Using the plus sign (+) or asterisk (*) to continue a numbered list using the %md interpreter results in bullet point entries instead.

Fixed Issues in Apache ZooKeeper

Review the list of ZooKeeper issues that are resolved in Cloudera Runtime 7.2.2.

CDPD-15702: Zookeeper clients fail to connect to a secure ZooKeeper if TLS certificate contains SAN entries that are not of type DNS or IP.

ZKHostnameVerifier now does not reject valid certificates with subjectAltNames anymore. This issue is now resolved.

Known Issues In Cloudera Runtime 7.2.2

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.2.2.

Known Issues in Apache Atlas

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

ATLAS-3921: Currently there is no migration path from AWS S3 version 1 to AWS S3 version 2.

None

CDPD-12668: Navigator Spark lineage can fail to render in Atlas

As part of content conversion from Navigator to Atlas, the conversion of some spark applications created a cyclic lineage reference in Atlas, which the Atlas UI fails to render. The cases occur when a Spark application uses data from a table and updates the same table.

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-11790: Simultaneous events on the Kafka topic queue can produce duplicate Atlas entities

In normal operation, Atlas receives metadata to create entities from multiple services on the same or separate Kafka topics. In some instances, such as for Spark jobs, metadata to create a table entity in Atlas is triggered from two separate messages: one for the Spark operation and a second for the table metadata from HMS. If the process metadata arrives before the table metadata, Atlas creates a temporary entity for any tables that are not already in Atlas and reconciles the temporary entity with the HMS metadata when the table metadata arrives.

However, in some cases such as when Spark SQL queries with the write.saveAsTable function, Atlas does not reconcile the temporary and final table metadata, resulting in two entities with the same qualified name and no lineage linking the table to the process entity.

This issue is not seen for other lineage queries from spark:

```
create table default.xx3 as select * from default.xx2
insert into yy2 select * from yy
insert overwrite table ww2 select * from ww1
```

Another case where this behavior may occur is when many REST API requests are sent at the same time.

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-11338: Cluster names with upper case letters may appear in lower case in some process names

Atlas records the cluster name as lower case in qualifiedNames for some process names. The result is that the cluster name may appear in lower case for some processes (insert overwrite table) while it appears in upper case for other queries (ctas) performed on the same cluster.

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-9095: Duplicate audits for renaming Hive tables

Renaming a Hive table results in duplicate ENTITY_UPDATE events in the corresponding Atlas entity audits, both for the table and for its columns.

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-7781: TLS certificates not validated on Firefox

Atlas is not checking for valid TLS certificates when the UI is opened in FireFox browsers.

None

CDPD-6675: Irregular qualifiedName format for Azure storage

The qualifiedName for hdfs_path entities created from Azure blob locations (ABFS) doesn't have the clusterName appended to it as do hdfs_path entities in other location types.

None

CDPD-5933 and 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 the 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 metadata collection from HMS. For example, if an ALTER TABLE operation in Spark changing a table name and is reported to Atlas before HMS has processed the change, Atlas may not show the correct lineage relationships to the altered table.

None

CDPD-4545: Searches for Qualified Names with "@" doesn't fetch the correct results

When searching Atlas qualifiedName values that include an "@" character (@), Atlas does not return the expected results or generate appropriate search suggestions.

Consider leaving out the portion of the search string that includes the @ sign, using the wildcard character * instead.

CDPD-3208: Table alias values are not found in search

When table names are changed, Atlas keeps the old name of the table in a list of aliases. These values are not included in the search index in this release, so after a table name is changed, searching on the old table name will not return the entity for the table.

None

CDPD-3160: Hive lineage missing for INSERT OVERWRITE queries

Lineage is not generated for Hive INSERT OVERWRITE queries on partitioned tables. Lineage is generated as expected for CTAS queries from partitioned tables.

None

CDPD-3125: Logging out of Atlas does not manage the external authentication

At this time, Atlas does not communicate a log-out event with the external authentication management, Apache Knox. When you log out of Atlas, you can still open the instance of Atlas from the same web browser without re-authentication.

To prevent access to Atlas after logging out, close all browser windows and exit the browser.

CDPD-1892: Ranking of top results in free-text search not intuitive

The Free-text search feature ranks results based on which attributes match the search criteria. The attribute ranking is evolving and therefore the choice of top results may not be intuitive in this release.

If you don't find what you need in the top 5 results, use the full results or refine the search.

CDPD-1884: Free text search in Atlas is case sensitive

The free text search bar in the top of the screen allows you to search across entity types and through all text attributes for all entities. The search shows the top 5 results that match the search terms at any place in the text (*term* logic). It also shows suggestions that match the search terms that begin with the term (term* logic). However, in this release, the search results are case-sensitive.

If you don't see the results you expect, repeat the search changing the case of the search terms.

CDPD-1823: Queries with ? wildcard return unexpected results

DSL queries in Advanced Search return incorrect results when the query text includes a question mark (?) wildcard character. This problem occurs in environments where trusted proxy for Knox is enabled, which is always the case for 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: IsUnique relationship attribute not honored

The Atlas model includes the ability to ensure that an attribute can be set to a specific value in only one relationship entity across the cluster metadata. For example, if you wanted to add metadata tags to relationships that you wanted to make sure were unique in the system, you could design the relationship attribute with the property "IsUnique" equal true. However, in this release, the IsUnique attribute is not enforced.

None

Known Issues in Data Analytics Studio

Learn about the known issues in Data Analytics Studio, the impact or changes to the functionality, and the workaround.

- You may not be able to add or delete columns or change the table schema after creating a new table using the upload table feature.

- For clusters secured using Knox, you see the HTTP 401: Forbidden error message when you click the DAS quick link from Cloudera Manager and are unable to log into DAS.

Workaround: The admin user will need to provide the DAS URL from the Knox proxy topology to the users needing access to DAS.

- The download logs feature may not return the YARN application logs on a Kerberized cluster. When you download the logs, the logs contain an error-reports.json file which states that no valid Kerberos tokens are available.

Workaround: An admin user with access to the machine can use the kinit command as a hive user with hive service user keytabs and trigger the download.

- The task logs for a particular task may not be available in the task swimlane. And the zip file generated by download logs artifact may not have task logs, but instead contain an error-reports.json file with the error log of the download failures.
- You may not see any data for a report for any new queries that you run. This can happen especially for the last one day's report.

Workaround:

1. Shut down the DAS Event Processor.
2. Run the following command from the Postgres server:

```
update das.report_scheduler_run_audit set status = 'FAILED' where status = 'READING';
```

3. Start the DAS Event Processor.

- On clusters secured with Knox proxy only: You might not be able to save the changes to the JDBC URL in the DAS UI to change the server interface (HS2 or LLAP) on which you are running your queries.
- You may be unable to upload tables or get an error while browsing files to upload tables in DAS on a cluster secured using Knox proxy.
- DAS does not parse semicolons (;) and double hyphens (--) in strings and comments.

For example, if you have a semicolon in query such as the following, the query might fail: `select * from properties where prop_value = "name1;name2";`

If a semicolon is present in a comment, then execute the query after removing the semicolon from the comment, or removing the comment altogether. For example:

```
select * from test; -- select * from test;
select * from test; /* comment; comment */
```

Queries with double hyphens (--) might also fail. For example:

```
select * from test where option = '--name';
```

- You might face UI issues on Google Chrome while using faceted search. We recommend you to use the latest version of Google Chrome (version 71.x or higher).
- Visual Explain for the same query shows different graphs on the **Compose** page and the **Query Details** page.
- While running some queries, if you restart HSI, the query execution is stopped. However, DAS does not reflect this change and the queries appear to be in the same state forever.
- After a fresh installation, when there is no data and you try to access the Reports tab, DAS displays an "HTTP 404 Not Found" error.
- Join count does not get updated for tables with partitioned columns.

Known Issues in Apache HBase

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

HBASE-24885: If an operator uses HBCK2 to invoke multiple `assigns` operations against one Region or happens to invoke HBCK2 `assigns` while HBase is re-assigning a Region, it is possible that the Region will be abnormally assigned. For example, unassigned, stuck in transition, and doubly-assigned.

Obtain a fix for this issue. Operators should definitely not schedule multiple assigns for a single Region at the same time, however there is still a potential race condition.

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

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

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

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

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

HDFS encryption with HBase

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

Workaround: N/A

AccessController postOperation problems in asynchronous operations

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

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

Workaround: N/A

Apache Issue: [HBASE-6992](#)

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

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

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

Apache Issue: N/A

Technical Service Bulletins

TSB 2021-453: Snapshot and cloned table corruption when original table is deleted

HBASE-25206 can cause data loss either through corrupting an existing hbase snapshot or destroying data that backs a clone of a previous snapshot.

Upstream JIRA

[HBASE-25206](#)

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-453: HBASE-25206 "snapshot and cloned table corruption when original table is deleted"](#).

TSB 2021-463: Snapshot and cloned table corruption when original table is deleted

The HDFS short-circuit setting `dfs.client.read.shortcircuit` is overwritten to disabled by `hbase-default.xml`. HDFS short-circuit reads bypass access to data in HDFS by using a domain socket (file) instead of a network socket. This alleviates the overhead of TCP to read data from HDFS which can have a meaningful improvement on HBase performance (as high as 30-40%).

Users can restore short-circuit reads by explicitly setting `dfs.client.read.shortcircuit` in HBase configuration via the configuration management tool for their product (e.g. Cloudera Manager or Ambari).

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-463: HBase Performance Issue](#).

TSB 2021-506: Active HBase MOB files can be removed

Actively used MOB files can be deleted by `MobFileCleanerChore` due to incorrect serialization of reference file names. This is causing data loss on MOB-enabled tables.

Upstream JIRA

- [HBASE-23723](#)
- [HBASE-25970](#)

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-506: Active HBase MOB files can be removed](#)

Known Issues in HDFS

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

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

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](#))
- More than two NameNodes ([HDFS-6440](#))
- 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.

OPSAPS-54299 Installing Hive on Tez and HMS in the incorrect order causes HiveServer failure

You need to install Hive on Tez and HMS in the correct order; otherwise, HiveServer fails.
 You need to install additional HiveServer roles to Hive on Tez, not the Hive service; otherwise, HiveServer fails.

Workaround: Follow instructions on [Installing Hive on Tez](#).

CDPD-15518: ACID tables you write using the Hive Warehouse Connector cannot be read from an Impala virtual warehouse.

Read the tables from a Hive virtual warehouse or using Impala queries in Data Hub.

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`)

ENGESC-2214: Hiveserver2 and HMS service logs are not deleted

Update Hive log4j configurations. Hive -> Configuration -> HiveServer2 Logging
 Advanced Configuration Snippet (Safety Valve) Hive Metastore -> Configuration ->
 Hive Metastore Server Logging Advanced Configuration Snippet (Safety Valve) Add
 the following to the configurations: `appender.DRFA.strategy.action.type=DELETE`
`appender.DRFA.strategy.action.basepath=${log.dir} appender.DRFA.strategy.action.maxdepth=1`
`appender.DRFA.strategy.action.PathConditions.glob=${log.file}.*`
`appender.DRFA.strategy.action.PathConditions.type=IfFileName`
`appender.DRFA.strategy.action.PathConditions.nestedConditions.type=IfAccumulatedFileCount`
`appender.DRFA.strategy.action.PathConditions.nestedConditions.exceeds=same value as`
`appender.DRFA.strategy.max`

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

CDPD-11890: Hive on Tez cannot run certain queries on tables stored in encryption zones

This problem occurs when the Hadoop Key Management Server (KMS) connection is SSL-encrypted and a self signed certificate is used. `SSLHandshakeException` might appear in Hive logs.

Use one of the workarounds:

- Install a self signed SSL certificate into `cacerts` file on all hosts.
- Copy `ssl-client.xml` to a directory that is available in all hosts. In Cloudera Manager, in Clusters Hive on Tez Configuration . In Hive Service Advanced Configuration Snippet for `hive-site.xml`, click +, and add the name `tez.aux.uris` and value `path-to-ssl-client.xml`.

Technical Service Bulletins**TSB 2021-459: Renaming managed (ACID) table shows empty records**

Renaming an ACID (managed) table using `ALTER TABLE <table name> RENAME` causes empty records in the table. Also, the location of the new table after renaming points to the location of the old table before renaming. This can cause correctness issues, for example:

```
create table abc (id int);
insert into abc values (1);
rename table abc to def; create table abc (id int); // should be
empty
insert into abc values (2);
select * from abc ; // returns 1 and 2, the new and the old re
sults
```

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-459: Renaming managed \(ACID\) table shows empty records](#)

TSB 2021-480/1: Hive produces incorrect query results when skipping a header in a binary file

In CDP, setting the table property `skip.header.line.count` to greater than 0 in a table stored in a binary format, such as Parquet, can cause incorrect query results. The skip header property is intended for use with Text files and typically used with CSV files. The issue is not present when you run the query on a Text file that sets the skip header property to 1 or greater.

Upstream JIRA

[Apache Jira: HIVE-24827](#)

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-480.1: Hive produces incorrect query results when skipping a header in a binary file](#)

TSB 2021-480/2: Hive ignores the property to skip a header or footer in a compressed file

In CDP, setting the table properties `skip.header.line.count` and `skip.footer.line.count` to greater than 0 in a table stored in a compressed format, such as bzip2, can cause incorrect results from `SELECT *` or `SELECT COUNT (*)` queries.

Upstream JIRA

[Apache Jira: HIVE-24224](#)

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-480.2: Hive ignores the property to skip a header or footer in a compressed file](#)

TSB 2021-482: Race condition in subdirectory delete/rename causes hive jobs to fail

Multiple threads try to perform a rename operation on s3. One of the threads fails to perform a rename operation, causing an error. Hive logs will report "HiveException: Error moving ..." and the log will contain an error line starting with " Exception when loading partition " -all paths listed with s3a:// prefixes.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-482: Race condition in subdirectory delete/rename causes Hive jobs to fail](#)

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.

Impact

Incorrect results

Action required

- **Hotfix request**

Request a hotfix from Cloudera Support.

- **Workaround**

Set `hive.vectorized.execution.mapjoin.native.fast.hashtable.enabled` to false. This might cause a performance degradation depending on the type of query and the system it is running on.

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 2021-518: Incorrect results returned when joining two tables with different bucketing versions

Incorrect results are returned when joining two tables with different bucketing versions, and with the following Hive configurations: set `hive.auto.convert.join` = false and set `mapreduce.job.reduces` = any custom value.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-518: Incorrect results returned when joining two tables with different bucketing versions](#)

TSB 2021-524: Intermittent data duplication if direct insert enabled

If direct insert is enabled, data is written directly to the final location with an attemptId. At the end of the insert operation, all data written before the final attempt should be deleted. However due to a bug in HIVE-21164, this does not happen.

Example: Data is written to the final location with attemptId=0, but this task fails. Hive tries the task again and writes data to the final location with attemptId=1. At the end of the insert, Hive should remove all the files with attemptId=0, but it does not.

Upstream JIRA

- [HIVE-21164](#)
- [HIVE-24322](#)

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-524: Intermittent data duplication if direct insert enabled](#)

TSB 2023-627: IN/OR predicate on binary column returns wrong result

An IN or an OR predicate involving a binary datatype column may produce wrong results. The OR predicate is converted to an IN due to the setting `hive.optimize.point.lookup` which is true by default. Only binary data types are affected by this issue. See <https://issues.apache.org/jira/browse/HIVE-26235> for example queries which may be affected.

Upstream JIRA

[HIVE-26235](#)

Knowledge article

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2023-627: IN/OR predicate on binary column returns wrong result](#)

Known Issues in Hue

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

Downloading Impala query results containing special characters in CSV format fails with ASCII codec error

In CDP, Hue is compatible with Python 2.7.x, but the Tablib library for Hue has been upgraded from 0.10.x to 0.14.x, which is generally used with the Python 3 release. If you try to download Impala query results having special characters in the result set in a CSV format, then the download may fail with the ASCII unicode decode error.

To fix this issue, downgrade the Tablib library to 0.12.x.

1. SSH into the Hue server host.
2. Change directory to the following:

```
cd /opt/cloudera/parcels/CDH-7.x/lib/
```

3. Back up the hue directory:

```
cp -R hue hue_original
```

4. Change to the hue directory:

```
cd hue
```

5. Install the Wheel package using pip:

```
./build/env/bin/pip install wheel
```

The Wheel package is used to avoid recompiling your software during every install.

6. Install the Python Setuptools package for Hue as follows:

```
./build/env/bin/pip install setuptools==44.1.0
```

7. Install Tablib version 0.12.1 as follows:

```
./build/env/bin/pip install tablib==0.12.1
```

8. Go to Cloudera Manager and restart the Hue service.

Impala SELECT table query fails with UTF-8 codec error

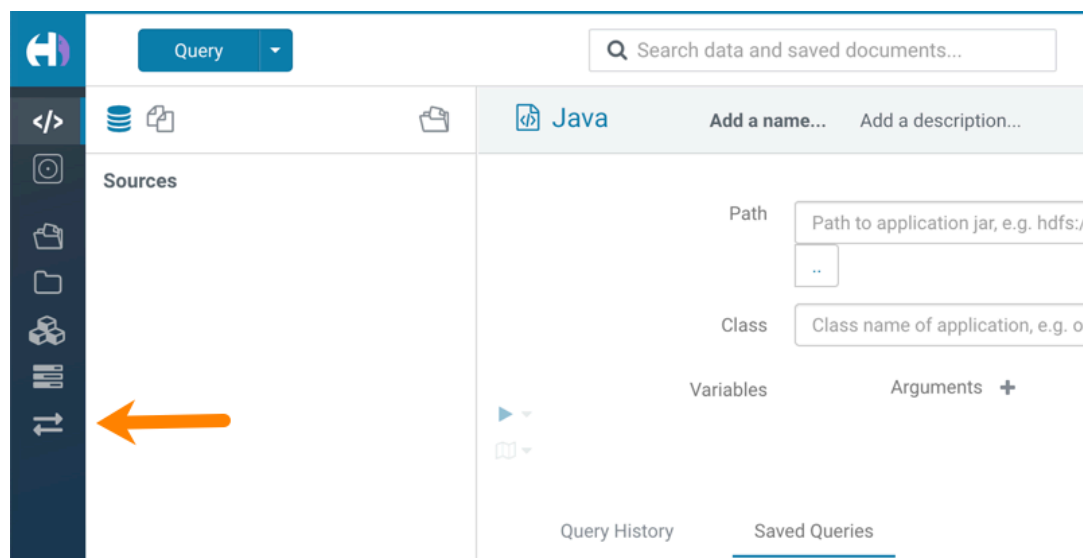
Hue cannot handle columns containing non-UTF8 data. As a result, you may see the following error while queying tables from the Impala editor in Hue: 'utf8' codec can't decode byte 0x91 in position 6: invalid start byte.

To resolve this issue, contact Cloudera Support to apply the following software patch: ENGESC-3457.

Hue Importer is not supported in the Data Engineering template

When you create a Data Hub cluster using the Data Engineering template, the Importer application is not supported in Hue.

Figure 1: Hue web UI showing Importer icon on the left assist panel



Hue Load Balancer role fails to start after upgrade to Cloudera Runtime 7 or you get the "BalancerMember worker hostname too long" error

You may see the following error message while starting the Hue Load Balancer:

```
BalancerMember worker hostname (xxx-xxxxxxx-xxxxxxxxxxx-xxxxxxx
.xxxxxx-xxxxxx-xxxxxx.example.site) too long.
```

Or, the Hue load balancer role fails to start after the upgrade, which prevents the Hue service from starting. If this failure occurs during cluster creation, cluster creation fails with the following error:

```
com.sequenceiq.cloudbreak.cm.ClouderaManagerOperationFailedException: Cluster template install failed: [Command [Start], with id [1234567890] failed:
Failed to start role., Command [Start], with id [1234567890] failed: Failed to start role., Command [Start], with id [1234567890] failed: Failed to start role.]
Unable to generate configuration for HUE_SERVER
Role failed to start due to error com.cloudera.cmf.service.config.ConfigGenException: Unable to generate config file hue.ini
```

Cloudera Manager displays this error when you create a Data Hub cluster using the Data Engineering template and the Hue Load Balancer worker node name has exceeded 64 characters. In a CDP Public Cloud deployment, the system automatically generates the Load Balancer worker node name through AWS or Azure.

For example, if you specify `cdp-123456-scalecluster` as the cluster name, CDP creates `cdp-123456-scalecluster-master2.repro-aw.a123-4a5b.example.site` as the worker node name.

Specify a shorter cluster name while creating a Data Hub cluster so that the final worker node name does not cross 64 characters.

For example, `cdp-123456-scale`.

Unsupported features

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 2021-487: Cloudera Hue is vulnerable to Cross-Site Scripting attacks

Multiple Cross-Site Scripting (XSS) vulnerabilities of Cloudera Hue have been found. They allow JavaScript code injection and execution in the application context.

- CVE-2021-29994 - The Add Description field in the Table schema browser does not sanitize user inputs as expected.
- CVE-2021-32480 - Default Home direct button in Filebrowser is also susceptible to XSS attack.
- CVE-2021-32481 - The Error snippet dialog of the Hue UI does not sanitize user inputs.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-487: Cloudera Hue is vulnerable to Cross-Site Scripting attacks \(CVE-2021-29994, CVE-2021-32480, CVE-2021-32481\)](#)

Known Issues in Apache Impala

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

Impala Virtual Warehouses might produce an error when querying transactional (ACID) tables

Problem: If you are querying transactional (ACID) tables with an Impala Virtual Warehouse and compaction is run on the compacting Hive Virtual Warehouse, the query might fail. The compacting process deletes files and the Impala Virtual Warehouse might not be aware of the deletion. Then when the Impala Virtual Warehouse attempts to read the deleted file, an error can occur. This situation occurs randomly.

Run the `INVALIDATE METADATA` statement on the transactional (ACID) table to refresh the metadata. This fixes the problem until the next compaction occurs.

HADOOP-15720: Queries stuck on failed HDFS calls and not timing out

In Impala 3.2 and higher, if the following error appears multiple times in a short duration while running a query, it would mean that the connection between the `impalad` and the HDFS NameNode is in a bad state.

```
"hdfsOpenFile() for <filename> at backend <hostname:port> failed
to finish before the <hdfs_operation_timeout_sec> second timeout
"
```

In Impala 3.1 and lower, the same issue would cause Impala to wait for a long time or not respond without showing the above error message.

Restart the `impalad`.

IMPALA-532: Impala should tolerate bad locale settings

If the `LC_*` environment variables specify an unsupported locale, Impala does not start.

Add `LC_ALL="C"` to the environment settings for both the Impala daemon and the Statestore daemon.

IMPALA-5605: Configuration to prevent crashes caused by thread resource limits

Impala could encounter a serious error due to resource usage under very high concurrency. The error message is similar to:

```
F0629 08:20:02.956413 29088 llvm-codegen.cc:111] LLVM hit fatal
error: Unable to allocate section memory!
terminate called after throwing an instance of 'boost::exception_
detail::clone_impl<boost::exception_detail::error_info_injector<
boost::thread_resource_error> >'
```

To prevent such errors, configure each host running an `impalad` daemon with the following settings:

```
echo 2000000 > /proc/sys/kernel/threads-max
echo 2000000 > /proc/sys/kernel/pid_max
echo 8000000 > /proc/sys/vm/max_map_count
```

Add the following lines in `/etc/security/limits.conf`:

```
impala soft nproc 262144
impala hard nproc 262144
```

IMPALA-635: Avro Scanner fails to parse some schemas

The default value in Avro schema must match type of first union type, e.g. if the default value is null, then the first type in the UNION must be "null".

Swap the order of the fields in the schema specification. For example, use ["null", "string"] instead of ["string", "null"]. Note that the files written with the problematic schema must be rewritten with the new schema because Avro files have embedded schemas.

IMPALA-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-9350: Ranger audit logs for applying column masking policies missing

Impala is not producing these logs.

None

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-1792: ImpalaODBC: Can not get the value in the SQLGetData(m-x th column) after the SQLBindCol(m th column)

If the ODBC SQLGetData is called on a series of columns, the function calls must follow the same order as the columns. For example, if data is fetched from column 2 then column 1, the SQLGetData call for column 1 returns NULL.

Fetch columns in the same order they are defined in the table.

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.

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.

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.

Technical Service Bulletins

TSB-2021-485: Impala returns fewer rows from parquet tables on S3

[IMPALA-10310](#) was an issue in Impala's Parquet page filtering code where the scanner did not reset state appropriately when transitioning from the first row group to subsequent row groups in a single split. This caused data from the subsequent row groups to be skipped incorrectly, leading to incorrect query results. This issue cannot occur when the Parquet page filtering is disabled by setting `PARQUET_READ_PAGE_INDEX=false`.

The issue is more likely to be encountered on S3/ADLS/ABFS/etc, because Spark is sometimes configured to write 128MB row groups and the `PARQUET_OBJECT_STORE_SPLIT_SIZE` is 256MB. This makes it more likely for Impala to process two row groups in a single split.

Parquet page filtering only works based on the min/max statistics, therefore the comparison operators it supports are "=", "<", ">", "<=", and ">=". These operators are impacted by this bug. Expressions such as "!=", 'LIKE' or the expressions including UDF do not use parquet page filtering.

The `PARQUET_OBJECT_STORE_SPLIT_SIZE` parameter is introduced in Impala 3.3 by [IMPALA-5843](#). This means that older versions of Impala do not have this issue.

Upstream JIRA

- [IMPALA-5843](#)
- [IMPALA-10310](#)

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-485: Impala returns fewer rows from parquet tables on S3](#)

TSB 2021-502: Impala logs the session / operation secret on most RPCs at INFO level

Impala logs contain the session / operation secret. With this information a person who has access to the Impala logs might be able to hijack other users' sessions. This means the attacker is able to execute statements for which they do not have the necessary privileges otherwise. Impala deployments where Apache Sentry or Apache Ranger authorization is enabled may be vulnerable to privilege escalation. Impala deployments where audit logging is enabled may be vulnerable to incorrect audit logging.

Restricting access to the Impala logs that expose secrets will reduce the risk of an attack. Additionally, restricting access to trusted users for the Impala deployment will also reduce the risk of an attack. Log redaction techniques can be used to redact secrets from the logs. For more information, see the *Cloudera Manager documentation*.

For log redaction, users can create a rule with a search pattern: `secret \ (string) [=:] .*` And the replacement could be for example: `secret=LOG-REDACTED`

Upstream JIRA

[IMPALA-10600](#)

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-502: Impala logs the session / operation secret on most RPCs at INFO level](#)

TSB 2021-479: Impala can return incomplete results through JDBC and ODBC clients in all CDP offerings

In CDP, we introduced a timeout on queries to Impala defaulting to 10 seconds. The timeout setting is called `FETCH_ROWS_TIMEOUT_MS`. Due to this setting, JDBC, ODBC, and Beeswax clients running Impala queries believe the data returned at 10 seconds is a complete dataset and present it as the final output. However, in cases where there are still results to return after this timeout has passed, when the driver closes the connection, based on the timeout, it results in a scenario where the query results are incomplete.

Upstream JIRA

[IMPALA-7561](#)

Knowledge article

For the latest update on this issue, see the corresponding Knowledge article: [TSB-2021 479: Impala can return incomplete results through JDBC and ODBC clients in all CDP offerings](#)

TSB 2022-543: Impala query with predicate on analytic function may produce incorrect results

Apache Impala may produce incorrect results for a query which has all of the following conditions:

- There are two or more analytic functions (for example, `row_number()`) in an inline view
- Some of the functions have partition-by expression while the others do not
- There is a predicate on the inline view's output expression corresponding to the analytic function

Upstream JIRA

[IMPALA-11030](#)

Knowledge article

For the latest update on this issue, see the corresponding Knowledge article: [TSB 2022-543: Impala query with predicate on analytic function may produce incorrect results](#)

Known Issues in Apache Kafka

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

Known Issues**OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners**

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

Workaround: SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You 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.

Topics created with the kafka-topics tool are only accessible by the user who created them when the deprecated --zookeeper option is used

By default all created topics are secured. However, when topic creation and deletion is done with the kafka-topics tool using the `--zookeeper` option, the tool talks directly to Zookeeper. Because security is the responsibility of ZooKeeper authorization and authentication, Kafka cannot prevent users from making ZooKeeper changes. As a result, if the `--zookeeper` option is used, only the user who created the topic will be able to carry out administrative actions on it. In this scenario Kafka will not have permissions to perform tasks on topics created this way.

Use kafka-topics with the `--bootstrap-server` option that does not require direct access to Zookeeper.

Certain Kafka command line tools require direct access to Zookeeper

The following command line tools talk directly to ZooKeeper and therefore are not secured via Kafka:

- kafka-reassign-partitions

None

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.

Custom Kerberos principal names cannot be used for kerberized ZooKeeper and Kafka instances

When using ZooKeeper authentication and a custom Kerberos principal, Kerberos-enabled Kafka does not start. You must disable ZooKeeper authentication for Kafka or use the default Kerberos principals for ZooKeeper and Kafka.

None

KAFKA-2561: Performance degradation when SSL Is enabled

In some configuration scenarios, significant performance degradation can occur when SSL is enabled. The impact varies depending on your CPU, JVM version, Kafka configuration, and message size. Consumers are typically more affected than producers.

Configure brokers and clients with ssl.secure.random.implementation = SHA1PRNG. It often reduces this degradation drastically, but its effect is CPU and JVM dependent.

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

OPSAPS-57907: The Kafka metric collector adapter generates high CPU load

If a large number of topic partitions are created on a cluster, the Cloudera Manager Agent can generate a high CPU load. This is caused by the Kafka metric collector adapter carrying out excessive regex matching.

None

OPSAPS-59031: Kafka cannot start if configuration is added to the Kafka Broker Advanced Configuration Snippet (Safety Valve) for ssl.properties

The Kafka Broker Advanced Configuration Snippet (Safety Valve) for ssl.properties configuration snippet does not correctly override configuration. As a result, Kafka may not start if TLS/SSL related configuration overrides are added to the this configuration snippet.

Use the Kafka Broker Advanced Configuration Snippet (Safety Valve) for kafka.properties configuration snippet instead to override SSL related properties.

OPSAPS-59032: Obsolete TLS protocols are enabled by default

TLS v1.0 and v1.1 are considered insecure and obsolete. Therefore, these protocols are disabled by default in Kafka 2.5. However for Cloudera Manager managed Kafka, these protocols are still enabled by default.

Add the following configuration to Kafka Broker Advanced Configuration Snippet (Safety Valve) for `kafka.properties`. This overrides the default and forces TLS v1.2 to be used.

```
ssl.protocol=TLSv1.2
ssl.enabled.protocols=TLSv1.2
```

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.
- While Kafka Connect is available as part of Runtime, it is currently not supported in CDP Public Cloud. NiFi is a proven solution for batch and real time data loading that complement Kafka's message broker capability. For more information, see [Creating your first Flow Management cluster](#).
- The Kafka default authorizer is not supported. This includes setting ACLs and all related APIs, broker functionality, and command-line tools.

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 HostsHosts 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.

OPSAPS-60331: If Cloudera Manager is configured to use Active Directory as a Kerberos KDC, and is also configured to use /etc/cloudera-scm-server/cmfs.keytab as the KDC admin credentials, you may encounter errors when generating Kerberos credentials.

In the Cloudera Manager Admin Console, run the "Administration > Security > Kerberos Credentials > Import KDC Account Manager Credentials" wizard. Remove /etc/cloudera-scm-server/cmfs.keytab on the Cloudera Manager server host.

Known Issues in Apache Knox

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

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.

Known Issues in Apache Kudu

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

Kudu supports only coarse-grain authorization. Kudu does not yet support integration with Atlas.

None

Kudu HMS Sync is disabled and is not yet supported

None

Known Issues in Apache Oozie

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

Oozie jobs fail (gracefully) on secure YARN clusters when JobHistory server is down

If the JobHistory server is down on a YARN (MRv2) cluster, Oozie attempts to submit a job, by default, three times. If the job fails, Oozie automatically puts the workflow in a SUSPEND state.

When the JobHistory server is running again, use the resume command to inform Oozie to continue the workflow from the point at which it left off.

CDPD-5340: The resourceManager property defined in an Oozie workflow might not work properly if the workflow is submitted through Knox proxy.

An Oozie workflow defined to use the resourceManager property might not work as expected in situations when the workflow is submitted through Knox proxy.

Define the jobTracker property with the same value as that of the resourceManager property.

Unsupported Feature

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

- Non-support for Pig action (CDPD-1070)
- Conditional coordinator input logic

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.

BUG-123856: Upgrade fails while configuring Oozie server.

None

Known Issues in Apache Phoenix

There are no known issues for Phoenix in Cloudera Runtime 7.2.2.

Known Issues in Apache Ranger

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

CDPD-3296: Audit files for Ranger plugin components do not appear immediately in S3 after cluster creation

For Ranger plugin components (Atlas, Hive, HBase, etc.), audit data is updated when the applicable audit file is rolled over. The default Ranger audit rollover time is 24 hours, so audit data appears 24 hours after cluster creation.

To see the audit logs in S3 before the default rollover time of 24 hours, use the following steps to override the default value in the Cloudera Manager safety valve for the applicable service.

1. On the Configuration tab in the applicable service, select Advanced under CATEGORY.
2. Click the + icon for the <service_name> Advanced Configuration Snippet (Safety Valve) for ranger-<service_name>-audit.xml property.
3. Enter the following property in the Name box:
xasecure.audit.destination.hdfs.file.rollover.sec.
4. Enter the desired rollover interval (in seconds) in the Value box. For example, if you specify 180, the audit log data is updated every 3 minutes.
5. Click Save Changes and restart the service.

CDPD-12644: Ranger Key Names cannot be reused with the Ranger KMS KTS service

Key names cannot be reused with the Ranger KMS KTS service. If the key name of a delete key is reused, the new key can be successfully created and used to create an encryption zone, but data cannot be written to that encryption zone.

Use only unique key names when creating keys.

Known Issues in Schema Registry

There are no known issues for Schema Registry in Cloudera Runtime 7.2.2.

Known Issues in Cloudera Search

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

Known Issues

TSB 2022-535: Ranger audit retention settings in Solr are not honored

The audits present in the ranger_audits collection in the Solr service of Data Lake do not get deleted based on the retention period set. The default retention period is 90 days.

This is caused by the incorrect order of processors in the configuration (solrconfig.xml) used by the ranger_audits collection.

Impact

- Older audits based on the retention period are not deleted that leads to:
 - An increase in space utilization under the Solr Data Directory. You can check the location being used for the “Solr Data Directory” setting in the Solr service configuration.
 - Increases the memory usage of Solr service
 - Backup-Restore of Solr collection would take longer

Action required

- **Workaround**
- Update Solr configuration and clean old audits

This is a two-step process. Note that this process must be completed for any Data Lake that has used an affected runtime version, even if the Data Lake has since been upgraded to a newer runtime.

1. Part 1: Manually correct the configurations to ensure that Solr retention value is enforced for new documents that are added.
2. Part 2: Manually cleanup documents that are older than 90 days.

- **Upgrade (recommended)**

Upgrade to runtime 7.2.11+

If the Data Lake is not upgraded, then the manual configuration Solr changes will be overwritten if the Data Lake is repaired or upgraded to runtime lower than 7.2.11. To ensure the changes are persisted, the Data Lake must be upgraded to runtime 7.2.11 or higher. See below for additional upgrade details. Note that the manual steps in Parts 1 and 2 above still need to be done even if the Data Lake is upgraded.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-535: Ranger audit retention settings in Solr are not honored](#).

Splitshard of 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.

None

OPSAPS-58059: Solr log rotation counts the number of retained log files daily instead of globally

With CDP 7.1.1, Search moved to Log4Jv2. This has affected Solr log rotation behavior in an unwanted way. With the default configuration, Solr log file names include a date and a running index, for example: solr-cmf-solr-SOLR_SERVER-solrserver-1.my.corporation.com.log.out.2020-08-31-9. The number of retained log files is configured in Cloudera Manager, however the configured number now applies for each day, instead of applying globally for all log files of the particular server.

Using Cloudera Manager, edit the Solr Server Logging Advanced Configuration Snippet (Safety Valve) property of your Solr service and add a new line containing: `appender.DRFA.filePattern=${log.dir}/${log.file}.%i`

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.

There is no workaround, you need to reindex collections.

CDH-82042: Solr service with no added collections causes the upgrade process to fail

Upgrade fails while performing the bootstrap collections step of the solr-upgrade.sh script with the error message:

```
Failed to execute command Bootstrap Solr Collections on service Solr
```

if there are no collections present in Solr.

If there are no collections added to it, remove the Solr service from your cluster before you start the upgrade.

CDH-34050: Collection Creation No Longer Supports Automatically Selecting A Configuration If Only One Exists

Before CDH 5.5.0, a collection could be created without specifying a configuration. If no -c value was specified, then:

- If there was only one configuration, that configuration was chosen.
- If the collection name matched a configuration name, that configuration was chosen.

Search now includes multiple built-in configurations. As a result, there is no longer a case in which only one configuration can be chosen by default.

Explicitly specify the collection configuration to use by passing -c <configName> to solrctl collect on --create.

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-19923: The quickstart.sh file does not validate ZooKeeper and the NameNode on some operating systems.

The quickstart.sh file uses the timeout function to determine if ZooKeeper and the NameNode are available. To ensure this check can be complete as intended, the quickstart.sh determines if the operating system on which the script is running supports timeout. If the script detects that the operating system does not support timeout, the script continues without checking if the NameNode and ZooKeeper are available. If your environment is configured properly or you are using an operating system that supports timeout, this issue does not apply.

This issue only occurs in some operating systems. If timeout is not available, the quickstart continues and final validation is always done by the MapReduce jobs and Solr commands that are run by the quickstart.

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. This must be done without Kerberos.

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.

Unsupported Features

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

- [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 Search. 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, eventually, may corrupt data collection for co-located core, or even shard and server level charts.

Known Issues in Apache Solr

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

Technical Service Bulletins

TSB 2021-495: CVE-2021-29943: Apache Solr Unprivileged users may be able to perform unauthorized read/write to collections

Using the `ConfigurableInternodeAuthHadoopPlugin` class as the authentication plugin with Ranger as the authorization module introduced a backdoor for unauthorized access to data. With this combination, when an authenticated user sends a query to a node, which does not have the data locally, the request will be forwarded in the name of the Solr service user and not in the name of the original requester. In this case, the authorization happens against the user named `solr` which may have almost full access. It may be the case that infra Solr customers were advised to switch back to `ConfigurableInternodeAuthHadoopPlugin`. Only these customers should be affected by this CVE.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-495: Apache Solr Unprivileged users may be able to perform unauthorized read/write to collections - CVE-2021-29943](#)

TSB 2021-497: CVE-2021-27905: Apache Solr SSRF vulnerability with the Replication handler

The Apache Solr ReplicationHandler (normally registered at `/replication` under a Solr core) has a `masterUrl` (also `leaderUrl` alias) parameter. The `masterUrl` parameter is used to designate another ReplicationHandler on another Solr core to replicate index data into the local core. To help prevent the CVE-2021-27905 SSRF vulnerability, Solr should check these parameters against a similar configuration used for the `shards` parameter.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-497: CVE-2021-27905: Apache Solr SSRF vulnerability with the Replication handler](#)

Known Issues in Apache Spark

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

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-217: HBase/Spark connectors are not supported

The *Apache HBase Spark Connector* (`hbase-connectors/spark`) and the *Apache Spark - Apache HBase Connector* (`shc`) are not supported in the initial CDP release.

None

CDPD-3038: Launching pyspark displays several HiveConf warning messages

When pyspark starts, several Hive configuration warning messages are displayed, similar to the following:

```
19/08/09 11:48:04 WARN conf.HiveConf: HiveConf of name hive.vect
orized.use.checked.expressions does not exist
19/08/09 11:48:04 WARN conf.HiveConf: HiveConf of name hive.te
z.cartesian-product.enabled does not exist
```

These errors can be safely ignored.

CDPD-2650: Spark cannot write ZSTD and LZ4 compressed Parquet to dynamically partitioned tables

Use a different compression algorithm.

CDPD-3783: Cannot create databases from Spark

Attempting to create a database using Spark results in an error similar to the following:

```
org.apache.spark.sql.AnalysisException:
  org.apache.hadoop.hive ql.metadata.HiveException: MetaException(message:Permission denied: user [sparkuser] does not have [ALL] privilege on [hdfs://ip-10-1-2-3.cloudera.site:8020/tmp/spark/warehouse/spark_database.db]);
```

Create the database using Hive or Impala, or specify the external data warehouse location in the create command. For example:

```
sql("create database spark_database location '/warehouse/tablespace/external/hive/spark_database.db' ")
```

Known Issues for Apache Sqoop

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

Using direct mode causes problems

Using direct mode has several drawbacks:

- Imports can cause intermittent an 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.

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.

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.

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.

OPSAPS-59553: SMM's bootstrap server config should be updated based on Kafka's listeners

SMM does not show any metrics for Kafka or Kafka Connect when multiple listeners are set in Kafka.

Workaround: SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You 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>
```

OPSAPS-59828: SMM cannot connect to Schema Registry when TLS is enabled

When TLS is enabled, SMM by default cannot properly connect to Schema Registry.

As a result, when viewing topics in the SMM Data Explorer with the deserializer key or value set to Avro, the following error messages are shown:

- Error deserializing key/value for partition [***PARTITION***] at offset [***OFFSET***]. If needed, please seek past the record to continue consumption.
- Failed to fetch value schema versions for topic : '[***TOPIC***]'.

In addition, the following certificate error will also be present in the SMM log:

- javax.net.ssl.SSLHandshakeException: PKIX path building failed:...

Workaround: Additional security properties must be set for SMM.

1. In Cloudera Manager, select the SMM service.
2. Go to Configuration.
3. Find and configure the SMM_JMX_OPTS property.

Add the following JVM SSL properties:

- Djavax.net.ssl.trustStore=[***SMM TRUSTSTORE LOCATION***]
- Djavax.net.ssl.trustStorePassword=[***PASSWORD***]

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

OPSAPS-57067: Yarn Service in Cloudera Manager reports stale configuration `yarn.cluster.scaling.recommendation.enable`.

This issue does not affect the functionality. Restarting Yarn service will fix this issue.

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.

OPSAPS-52066: Stacks under Logs Directory for Hadoop daemons are not accessible from Knox Gateway.

Stacks under the Logs directory for Hadoop daemons, such as NameNode, DataNode, ResourceManager, NodeManager, and JobHistoryServer are not accessible from Knox Gateway.

Administrators can SSH directly to the Hadoop Daemon machine to collect stacks under the Logs directory.

CDPD-2936: Application logs are not accessible in WebUI2 or Cloudera Manager

Running Containers Logs from NodeManager local directory cannot be accessed either in Cloudera Manager or in WebUI2 due to log aggregation.

Use the YARN log CLI to access application logs. For example:

```
yarn logs -applicationId <APPLICATION ID>
```

Apache Issue: [YARN-9725](#)

OPSAPS-50291: Environment variables HADOOP_HOME, PATH, LANG, and TZ are not getting whitelisted

It is possible to whitelist the environment variables `HADOOP_HOME`, `PATH`, `LANG`, and `TZ`, but the container launch environments do not have these variables set up automatically.

You can manually add the required environment variables to the whitelist using Cloudera Manager.

1. In Cloudera Manager, select the YARN service.
2. Click the Configuration tab.
3. Search for Containers Environment Variable Whitelist.
4. Add the environment variables (`HADOOP_HOME`, `PATH`, `LANG`, `TZ`) which are required to the list.
5. Click Save Changes.
6. Restart all NodeManagers.
7. Check the YARN aggregated logs to ensure that newly whitelisted environment variables are set up for container launch.

COMPX-3181: Application logs does not work for AZURE and AWS cluster

Yarn Application Log Aggregation will fail for any YARN job (MR, Tez, Spark, etc) which do not use cloud storage, or use a cloud storage location other than the one configured for YARN logs (`yarn.nodemanager.remote-app-log-dir`).

Configure the following:

- For MapReduce job, set `mapreduce.job.hdfs-servers` in the `mapred-site.xml` file with all filesystems required for the job including the one set in `yarn.nodemanager.remote-app-log-dir` such as `hdfs://nn1/`, `hdfs://nn2/`.
- For Spark job, set the job level with all filesystems required for the job including the one set in `yarn.nodemanager.remote-app-log-dir` such as `hdfs://nn1/`, `hdfs://nn2/` in `spark.yarn.access.hadoopFileSystems` and pass it through the `--config` option in `spark-submit`.
- For jobs submitted using the `hadoop` command, place a separate `core-site.xml` file with `fs.defaultFS` set to the filesystem set in `yarn.nodemanager.remote-app-log-dir` in a path. Add that directory path in `--config` when executing the `hadoop` command.

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 Resource

When YARN High Availability is enabled there are multiple Resource Managers. 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.

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.

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 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.

Unsupported Features

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

- GPU support for Docker
- Hadoop Pipes
- Fair Scheduler
- Application Timeline Server (ATS 2 and ATS 1.5)
- Container Resizing
- Distributed or Centralized Allocation of Opportunistic Containers
- Distributed Scheduling
- Native Services
- Pluggable Scheduler Configuration
- Queue Priority Support
- Reservation REST APIs
- Resource Estimator Service
- Resource Profiles
- (non-Zookeeper) ResourceManager State Store
- Shared Cache
- YARN Federation
- Rolling Log Aggregation
- Docker on YARN (DockerContainerExecutor) on Data Hub clusters
- Dynamic Resource Pools

Known Issues in Apache Zeppelin

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

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.

Known Issues in Apache ZooKeeper

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

Zookeeper-client does not use ZooKeeper TLS/SSL automatically

The command-line tool 'zookeeper-client' is installed to all Cloudera Nodes and it can be used to start the default Java command line ZooKeeper client. However even when ZooKeeper TLS/SSL is enabled, the zookeeper-client command connects to localhost:2181, without using TLS/SSL.

Manually configure the 2182 port, when zookeeper-client connects to a ZooKeeper cluster. The following is an example of connecting to a specific three-node ZooKeeper cluster using TLS/SSL:

```
CLIENT_JVMFLAGS="-Dzookeeper.clientCnxnSocket=org.apache.zoo
keeper.ClientCnxnSocketNetty -Dzookeeper.ssl.keyStore.locati
on=<PATH TO YOUR CONFIGURED KEYSTORE> -Dzookeeper.ssl.keyStor
e.password=<THE PASSWORD YOU CONFIGURED FOR THE KEYSTORE> -
Dzookeeper.ssl.trustStore.location=<PATH TO YOUR CONFIGURED
TRUSTSTORE> -Dzookeeper.ssl.trustStore.password=<THE PASSWORD
YOU CONFIGURED FOR THE TRUSTSTORE> -Dzookeeper.client.secu
re=true" zookeeper-client -server <YOUR.ZOOKEEPER.SERVER-1>:218
2,<YOUR.ZOOKEEPER.SERVER-2>:2182,<YOUR.ZOOKEEPER.SERVER-3>:2182
```

Behavioral Changes In Cloudera Runtime 7.2.2

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 7.2.2.

Behavioral Changes in Apache Phoenix

Learn about the change in certain functionality of Apache Phoenix that has resulted in a change in behavior from the previously released version to this version of Cloudera Runtime.

Summary:

As a consequence of including the Apache Phoenix fix PHOENIX-5213, some libraries that are in the old shaded phoenix-client JARs are included in the new JAR without relocation. These libraries are shaded in the new JAR. For example, com.google.protobuf package is shaded in the new JAR.

New behavior:

If you use any of these JARs in your current deployment, you must explicitly add the newly shaded dependencies to your applications.

Behavioral Changes in Ranger

Learn about the change in certain functionality of Cloudera Ranger that has resulted in a change in behavior from the previously released version to this version of Cloudera Runtime.

Cloudera ID

CDPD-4141

Apache component

Ranger

Apache JIRA:

N/A

Summary:

Default location for ranger audit data collected by Solr changed to local file system from HDFS

Details:**Previous behavior:**

The default value of the index storage in HDFS is /solr-infra.

New behavior:

The default value of the index storage in the local file system is /var/lib/solr-infra.

Behavioral Changes in Apache Hadoop YARN

Learn about the change in certain functionality of Hadoop YARN 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 `yarn.nodemanager.resource.memory-mb` and the `yarn.scheduler.maximum-allocation-mb` changed.

Previous behavior:

Default value of the `yarn.nodemanager.resource.memory-mb` and the `yarn.scheduler.maximum-allocation-mb` properties used to be 3GB.

New behavior:

Default value of the `yarn.nodemanager.resource.memory-mb` and the `yarn.scheduler.maximum-allocation-mb` properties are 9.64GB.

Deprecation Notices In Cloudera Runtime 7.2.2

Certain features and functionalities have been removed or deprecated in Cloudera Runtime 7.2.2. 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.

Removed Components and Product Capabilities

No components are deprecated or removed in this Cloudera Runtime release.

Please contact Cloudera Support or your Cloudera Account Team if you have any questions.

Deprecation notices in Apache Kudu

Certain features and functionality in Kudu are deprecated or removed in Cloudera Runtime 7.2.2. You must review these changes along with the information about the features in Kudu that will be removed or deprecated in a future release.

- The Flume sink has been migrated to the Apache Flume project and removed from Kudu. Users depending on the Flume integration can use the old kudu-flume jars or migrate to the Flume jars containing the Kudu sink.
- Support for Apache Sentry authorization has been deprecated and may be removed in the next release. Users depending on the Sentry integration should migrate to the Apache Ranger integration for authorization.
- Support for Python 2 has been deprecated and may be removed in the next release.
- Support for CentOS/RHEL 6, Debian 8, Ubuntu 14 has been deprecated and may be removed in the next release.

Deprecation Notices for Apache Kafka

Certain features and functionality in Kafka are deprecated or removed in Cloudera Runtime 7.2.2. You must review these changes along with the information about the features in Kafka that will be removed or deprecated in a future release.

Deprecated

kafka-preferred-replica-election

The `kafka-preferred-replica-election.sh` command line tool has been deprecated in upstream Apache Kafka 2.4.0. Its alternative in CDP, `kafka-preferred.replica-election`, is also deprecated.

--zookeeper

The `--zookeeper` option has been deprecated for all Kafka command line tools except for `kafka-reassign-partitions`. Cloudera recommends that you use the `--bootstrap-server` option instead.

Deprecation Notices in Apache HBase

Certain features and functionality in HBase are deprecated or removed in Cloudera Runtime 7.2.2. You must review these changes along with the information about the features in HBase that will be removed or deprecated in a future release.

Use this list to understand some of the deprecated items and incompatibilities if you are upgrading from HDP 2.x or CDH 5.x to CDP.

Known Incompatibilities when Upgrading from CDH and HDP

Cloudera Runtime uses Apache HBase 2.x.x whereas CDH 5.x and HDP 2.x uses Apache HBase 1.x.



Important: Some APIs that are listed as deprecated, but these APIs do not block your upgrade. You must stop using the deprecated APIs in your existing applications after upgrade, and not use these APIs in new development.

List of Major Changes

- HBASE-16189 and HBASE-18945: You cannot open the Cloudera Runtime HFiles in CDH or HDP.
- HBASE-18240: Changed the ReplicationEndpoint Interface.
- The Dynamic Jars Directory property `hbase.dynamic.jars.dir` is disabled by default. If you want to enable dynamic classloading, you can use the `hbase.dynamic.jars.dir` property in Cloudera Manager to change the default `${hbase.rootdir}/lib` directory to some other location, preferably a location on HDFS. This property is flagged by

Cloudera Manager as deprecated when you upgrade to CDP because the property is incompatible with HBase on cloud deployments. If you are using HBase with HDFS storage, you can ignore this warning, and keep using the `hbase.use.dynamic.jars` feature.

Co-processor API changes

- HBASE-16769: Deprecated Protocol Buffers references from `MasterObserver` and `RegionServerObserver`.
- HBASE-17312: [JDK8] Use default method for `Observer` Coprocessors. The interface classes of `BaseMasterAndRegionObserver`, `BaseMasterObserver`, `BaseRegionObserver`, `BaseRegionServerObserver` and `BaseWALObserver` uses JDK8's 'default' keyword to provide empty and no-op implementations.
- Interface `HTableInterface` introduces following changes to the methods listed below:

[#] interface `CoprocessorEnvironment`

Change	Result
Abstract method <code>getTable (TableName)</code> has been removed.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>getTable (TableName, ExecutorService)</code> has been removed.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.

- Public Audience

The following tables describes the coprocessor changes:

[#] class `CoprocessorRpcChannel` (1)

Change	Result
This class has become interface.	A client program may be interrupted by <code>IncompatibleClassChangeError</code> or <code>InstantiationException</code> exception depending on the usage of this class.

Class `CoprocessorHost<E>`

Classes that were Audience Private but were removed:

Change	Result
Type of field <code>coprocessors</code> has been changed from <code>java.util.SortedSet<E></code> to <code>org.apache.hadoop.hbase.util.SortedList<E></code> .	A client program may be interrupted by <code>NoSuchFieldError</code> exception.

MasterObserver changes

The following changes are introduced to the `MasterObserver` interface:

[#] interface `MasterObserver` (14)

Change	Result
Abstract method <code>voidpostCloneSnapshot (ObserverContext<MasterCoprocessorEnvironment>, HBaseProtos.SnapshotDescription, HTableDescriptor)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>voidpostCreateTable (ObserverContext<MasterCoprocessorEnvironment>, HTableDescriptor, HRegionInfo[])</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>voidpostDeleteSnapshot (ObserverContext<MasterCoprocessorEnvironment>, HBaseProtos.SnapshotDescription)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>voidpostGetTableDescriptors (ObserverContext<MasterCoprocessorEnvironment>, List<HTableDescriptor>)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.

Abstract method voidpostModifyTable (ObserverContext<MasterCoproprocessorEnvironment>, TableName, HTableDescriptor) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpostRestoreSnapshot (ObserverContext<MasterCoproprocessorEnvironment>, HBaseProtos.SnapshotDescription, HTableDescriptor) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpostSnapshot (ObserverContext<MasterCoproprocessorEnvironment>, HBaseProtos.SnapshotDescription, HTableDescriptor) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpreCloneSnapshot (ObserverContext<MasterCoproprocessorEnvironment>, HBaseProtos.SnapshotDescription, HTableDescriptor) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpreCreateTable (ObserverContext<MasterCoproprocessorEnvironment>, HTableDescriptor, HRegionInfo[]) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpreDeleteSnapshot (ObserverContext<MasterCoproprocessorEnvironment>, HBaseProtos.SnapshotDescription) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpreGetTableDescriptors (ObserverContext<MasterCoproprocessorEnvironment>, List<TableName>, List<HTableDescriptor>) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpreModifyTable (ObserverContext<MasterCoproprocessorEnvironment>, TableName, HTableDescriptor) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpreRestoreSnapshot (ObserverContext<MasterCoproprocessorEnvironment>, HBaseProtos.SnapshotDescription, HTableDescriptor) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpreSnapshot (ObserverContext<MasterCoproprocessorEnvironment>, HBaseProtos.SnapshotDescription, HTableDescriptor) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.

RegionObserver interface changes

The following changes are introduced to the RegionObserver interface.

[#] interface RegionObserver (13)

Change	Result
Abstract method voidpostCloseRegionOperation (ObserverContext<RegionCoproprocessorEnvironment>, HRegion.Operation) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpostCompactSelection (ObserverContext<RegionCoproprocessorEnvironment>, Store, ImmutableList<StoreFile>) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpostCompactSelection (ObserverContext<RegionCoproprocessorEnvironment>, Store, ImmutableList<StoreFile>, CompactionRequest) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpostGetClosestRowBefore (ObserverContext<RegionCoproprocessorEnvironment>, byte[], byte[], Result) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method DeleteTrackerpostInstantiateDeleteTracker (ObserverContext<RegionCoproprocessorEnvironment>, DeleteTracker) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpostSplit (ObserverContext<RegionCoproprocessorEnvironment>, HRegion, HRegion) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.

Abstract method voidpostStartRegionOperation (ObserverContext<RegionCoproprocessorEnvironment>, HRegionInfo, HLogKey, WALEdit) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method StoreFile.ReaderpostStoreFileReaderOpen (ObserverContext<RegionCoproprocessorEnvironment>, FileSystem, Path, FSDataInputStreamWrapper, long, CacheConfig, Reference, StoreFile.Reader) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method voidpostWALRestore (ObserverContext<RegionCoproprocessorEnvironment>, HRegionInfo, HLogKey, WALEdit) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method InternalScannerpreFlushScannerOpen (ObserverContext<RegionCoproprocessorEnvironment>, Store, KeyValueScanner, InternalScanner) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method voidpreGetClosestRowBefore (ObserverContext<RegionCoproprocessorEnvironment>, byte[], byte[], Result) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method StoreFile.ReaderpreStoreFileReaderOpen (ObserverContext<RegionCoproprocessorEnvironment>, FileSystem, Path, FSDataInputStreamWrapper, long, CacheConfig, Reference, StoreFile.Reader) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method voidpreWALRestore (ObserverContext<RegionCoproprocessorEnvironment>, HRegionInfo, HLogKey, WALEdit) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.

WALObserver interface changes

The following changes are introduced to the WALObserver interface:

[#] interface WALObserver

Change	Result
Abstract method voidpostWALWrite (ObserverContext<WALCoproprocessorEnvironment>, HRegionInfo, HLogKey, WALEdit) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method booleanpreWALWrite (ObserverContext<WALCoproprocessorEnvironment>, HRegionInfo, HLogKey, WALEdit) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.

Scheduler changes

Following methods are now changed to abstract:

[#]class RpcScheduler (1)

Change	Result
Abstract method void dispatch (CallRunner) has been removed from this class.	A client program may be interrupted by NoSuchMethodError exception.

[#] RpcScheduler.dispatch (CallRunner p1) [abstract] : void 1

org/apache/hadoop/hbase/ipc/RpcScheduler.dispatch:(Lorg/apache/hadoop/hbase/ipc/CallRunner;)V

Change	Result
Return value type has been changed from void to boolean.	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

The following abstract methods have been removed:

[#]interface PriorityFunction (2)

Change	Result
Abstract method <code>longgetDeadline (RPCProtos.RequestHeader, Message)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>int getPriority (RPCProtos.RequestHeader, Message)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.

Server API changes

[#] class `RpcServer` (12)

Change	Result
Type of field <code>CurCall</code> has been changed from <code>java.lang.ThreadLocal<RpcServer.Call></code> to <code>java.lang.ThreadLocal<RpcCall></code> .	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Abstract method <code>int getNumOpenConnections ()</code> has been added to this class.	This class became abstract and a client program may be interrupted by <code>InstantiationException</code> exception.
Field <code>callQueueSize</code> of type <code>org.apache.hadoop.hbase.util.Counter</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>connectionList</code> of type <code>java.util.List<RpcServer.Connection></code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>maxIdleTime</code> of type <code>int</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>numConnections</code> of type <code>int</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>port</code> of type <code>int</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>purgeTimeout</code> of type <code>long</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>responder</code> of type <code>RpcServer.Responder</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>socketSendBufferSize</code> of type <code>int</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>thresholdIdleConnections</code> of type <code>int</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.

Following abstract methods are removed:

Change	Result
Abstract method <code>Pair<Message,CellScanner>call (BlockingService, Descriptors.MethodDescriptor, Message, CellScanner, long, MonitoredRPCHandler)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.

Replication and WAL changes

HBASE-18733: `WALKey` has been purged completely. Following are the changes to the `WALKey`:

[#] class `WALKey` (8)

Change	Result
Access level of field <code>clusterIds</code> has been changed from protected to private.	A client program may be interrupted by <code>IllegalAccessError</code> exception.
Access level of field <code>compressionContext</code> has been changed from protected to private.	A client program may be interrupted by <code>IllegalAccessError</code> exception.
Access level of field <code>encodedRegionName</code> has been changed from protected to private.	A client program may be interrupted by <code>IllegalAccessError</code> exception.
Access level of field <code>tablename</code> has been changed from protected to private.	A client program may be interrupted by <code>IllegalAccessError</code> exception.

Access level of field writeTime has been changed from protected to private.	A client program may be interrupted by IllegalAccessException exception.
---	--

Following fields have been removed:

Change	Result
Field LOG of type org.apache.commons.logging.Log has been removed from this class.	A client program may be interrupted by NoSuchFieldError exception.
Field VERSION of type WALKey.Version has been removed from this class.	A client program may be interrupted by NoSuchFieldError exception.
Field logSeqNum of type long has been removed from this class.	A client program may be interrupted by NoSuchFieldError exception.

Admin Interface API changes

You cannot administer a CDP Runtime Data Hub cluster using a client that includes RelicationAdmin, ACC, Thrift and REST usage of Admin ops. Methods returning protobufs have been changed to return POJOs instead. Returns have changed from void to Future for async methods. HBASE-18106 - Admin.listProcedures and Admin.listLocks were renamed to getProcedures and getLocks. MapReduce makes use of Admin doing following admin.getClusterStatus() to calculate Splits.

- Thrift usage of Admin API:

```
compact(ByteBuffer) createTable(ByteBuffer, List<ColumnDescriptor>) deleteTable(ByteBuffer) disableTable(ByteBuffer) enableTable(ByteBuffer) getTableNames() majorCompact(ByteBuffer)
```

- REST usage of Admin API:

```
hbase-rest org.apache.hadoop.hbase.rest RootResource getTableList() TableName[] tableNames = servlet.getAdmin().listTableNames(); SchemaResource delete(UriInfo) Admin admin = servlet.getAdmin(); update(TableSchemaModel, boolean, UriInfo) Admin admin = servlet.getAdmin(); StorageClusterStatusResource get(UriInfo) ClusterStatus status = servlet.getAdmin().getClusterStatus(); StorageClusterVersionResource get(UriInfo) model.setVersion(servlet.getAdmin().getClusterStatus().getHBaseVersion()); TableResource exists() return servlet.getAdmin().tableExists(TableName.valueOf(table));
```

[#] interface Admin (9)

Following are the changes to the Admin interface:

Change	Result
Abstract method createTableAsync (HTableDescriptor, byte[] []) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method disableTableAsync (TableName) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method enableTableAsync (TableName) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method getCompactionState (TableName) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method getCompactionStateForRegion (byte[]) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method isSnapshotFinished (HBaseProtos.SnapshotDescription) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method snapshot (String, TableName, HBaseProtos.SnapshotDescription.Type) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.

Abstract method snapshot (HBaseProtos.SnapshotDescription) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method takeSnapshotAsync (HBaseProtos.SnapshotDescription) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.

[#] Admin.createTableAsync (HTableDescriptor p1, byte[][] p2) [abstract] : void 1

org/apache/hadoop/hbase/client/Admin.createTableAsync:(Lorg/apache/hadoop/hbase/HTableDescriptor;[[B)V

Change	Result
Return value type has been changed from void to java.util.concurrent.Future<java.lang.Void>.	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

[#] Admin.disableTableAsync (TableName p1) [abstract] : void 1

org/apache/hadoop/hbase/client/Admin.disableTableAsync:(Lorg/apache/hadoop/hbase/TableName;)V

Change	Result
Return value type has been changed from void to java.util.concurrent.Future<java.lang.Void>.	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

Admin.enableTableAsync (TableName p1) [abstract] : void 1

org/apache/hadoop/hbase/client/Admin.enableTableAsync:(Lorg/apache/hadoop/hbase/TableName;)V

Change	Result
Return value type has been changed from void to java.util.concurrent.Future<java.lang.Void>.	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

Admin.enableTableAsync (TableName p1) [abstract] : void 1

org/apache/hadoop/hbase/client/Admin.getCompactionState:(Lorg/apache/hadoop/hbase/TableName;)Lorg/apache/hadoop/hbase/protobuf/generated/AdminProtos\$GetRegionInfoResponse\$CompactionState;

Change	Result
Return value type has been changed from org.apache.hadoop.hbase.protobuf.generated.AdminProtos.GetRegionInfoResponse.CompactionState to CompactionState.	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

[#] Admin.getCompactionStateForRegion (byte[] p1) [abstract] : AdminProtos.GetRegionInfoResponse.CompactionState 1

org/apache/hadoop/hbase/client/Admin.getCompactionStateForRegion:([B)Lorg/apache/hadoop/hbase/protobuf/generated/AdminProtos\$GetRegionInfoResponse\$CompactionState;

Change	Result
Return value type has been changed from org.apache.hadoop.hbase.protobuf.generated.AdminProtos.GetRegionInfoResponse.CompactionState to CompactionState.	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

HTableDescriptor and HColumnDescriptor changes

HTableDescriptor and HColumnDescriptor has become interfaces and you can create it through Builders. HCD has become CFD. It no longer implements writable interface. package org.apache.hadoop.hbase.

[#] class HColumnDescriptor (1)

Change	Result
--------	--------

Removed super-interface <code>org.apache.hadoop.io.WritableComparable<HColumnDescriptor></code> .	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
---	--

class `HTableDescriptor` (3)

Change	Result
Removed super-interface <code>org.apache.hadoop.io.WritableComparable<HTableDescriptor></code> .	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Field <code>META_TABLEDESC</code> of type <code>HTableDescriptor</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.

[#] `HTableDescriptor.getColumnFamilies () : HColumnDescriptor[]` (1)

`org/apache/hadoop/hbase/HTableDescriptor.getColumnFamilies():[Lorg/apache/hadoop/hbase/HColumnDescriptor;`

[#] class `HColumnDescriptor` (1)

Change	Result
Return value type has been changed from <code>HColumnDescriptor[]</code> to <code>ClientColumnFamilyDescriptor[]</code> .	This method has been removed because the return type is part of the method signature. A client program may be interrupted by <code>NoSuchMethodError</code> exception.

[#] interface `Table` (4)

Change	Result
Abstract method <code>batch (List<?>)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>batchCallback (List<?>, Batch.Callback<R>)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>getWriteBufferSize ()</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>setWriteBufferSize (long)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.

Deprecated buffer methods

- `LockTimeoutException` and `OperationConflictException` classes have been removed.

class `OperationConflictException` (1)

Result	Result
This class has been removed.	A client program may be interrupted by <code>NoClassDefFoundError</code> exception.

class `LockTimeoutException` (1)

Change Result This class has been removed. A client program may be interrupted by `NoClassDefFoundError` exception.

Filter API changes

Following methods have been removed: package `org.apache.hadoop.hbase.filter`

[#] class `Filter` (2)

Result	Result
Abstract method <code>getNextKeyHint (KeyValue)</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.

Abstract method transform (KeyValue) has been removed from this class.	A client program may be interrupted by NoSuchMethodError exception.
--	---

- HBASE-12296: Filters should work with ByteBufferedCell.
- HConnection is removed in Cloudera Runtime.
- RegionLoad and ServerLoad internally moved to shaded Protocol Buffers.

[#] class RegionLoad (1)

Result	Result
Type of field regionLoadPB has been changed from protobuf.generated.ClusterStatusProtos.RegionLoad to shaded.protobuf.generated.ClusterStatusProtos.RegionLoad.	A client program may be interrupted by NoSuchFieldError exception.

[#] interface AccessControlConstants (3)

Result	Result
Field OP_ATTRIBUTE_ACL_STRATEGY of type java.lang.String has been removed from this interface.	A client program may be interrupted by NoSuchFieldError exception.
Field OP_ATTRIBUTE_ACL_STRATEGY_CELL_FIRST of type byte[] has been removed from this interface.	A client program may be interrupted by NoSuchFieldError exception.
Field OP_ATTRIBUTE_ACL_STRATEGY_DEFAULT of type byte[] has been removed from this interface.	A client program may be interrupted by NoSuchFieldError exception.

[#] ServerLoad.getNumberOfRequests () : int 1

org/apache/hadoop/hbase/ServerLoad.getNumberOfRequests():I

Result	Result
Return value type has been changed from int to long.	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

[#] ServerLoad.getNumberOfRequests () : int 1

org/apache/hadoop/hbase/ServerLoad.getReadRequestsCount():I

Result	Result
Return value type has been changed from int to long.	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

[#] ServerLoad.getTotalNumberOfRequests () : int 1

org/apache/hadoop/hbase/ServerLoad.getTotalNumberOfRequests():I

Result	Result
Return value type has been changed from int to long.	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

[#]ServerLoad.getWriteRequestsCount () : int 1

org/apache/hadoop/hbase/ServerLoad.getWriteRequestsCount():I

Result	Result
Return value type has been changed from int to long.	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

[#]class HConstants (6)

Result	Result
Field DEFAULT_HBASE_CONFIG_READ_ZOOKEEPER_CONFIG of type boolean has been removed from this class.	A client program may be interrupted by NoSuchFieldError exception.
Field HBASE_CONFIG_READ_ZOOKEEPER_CONFIG of type java.lang.String has been removed from this class.	A client program may be interrupted by NoSuchFieldError exception.
Field REPLICATION_ENABLE_DEFAULT of type boolean has been removed from this class.	A client program may be interrupted by NoSuchFieldError exception.
Field REPLICATION_ENABLE_KEY of type java.lang.String has been removed from this class.	A client program may be interrupted by NoSuchFieldError exception.
Field ZOOKEEPER_CONFIG_NAME of type java.lang.String has been removed from this class.	A client program may be interrupted by NoSuchFieldError exception.
Field ZOOKEEPER_USEMULTI of type java.lang.String has been removed from this class.	A client program may be interrupted by NoSuchFieldError exception.

HBASE-18732: [compat 1-2] HBASE-14047 removed Cell methods without deprecation cycle.

[#]interface Cell 5

Result	Result
Abstract method getFamily () has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method getMvccVersion () has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method getQualifier () has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method getRow () has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.
Abstract method getValue () has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.

HBASE-18795:Expose KeyValue.getBuffer() for tests alone. Allows KV#getBuffer in tests only that was deprecated previously.

Region scanner changes

[#]interface RegionScanner (1)

Result	Result
Abstract method boolean nextRow (List<Cell>, int) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.

StoreFile changes

[#] class StoreFile (1)

Result	Result
This class became interface.	A client program may be interrupted by IncompatibleClassChangeError or InstantiationException exception dependent on the usage of this class.

MapReduce changes

HFile*Format has been removed.

ClusterStatus changes

[#] ClusterStatus.getRegionsInTransition () : Map<String,RegionState> 1

org/apache/hadoop/hbase/ClusterStatus.getRegionsInTransition:()Ljava/util/Map;

Result	Result
Return value type has been changed from java.util.Map<java.lang.String,master.RegionState> to java.util.List<master.RegionState>.	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

Other changes in ClusterStatus include removal of convert methods that were no longer necessary after purge of Protocol Buffers from API.

Purge of Protocol Buffers from API

Protocol Buffers (PB) has been deprecated in APIs.

[#] HBaseSnapshotException.getSnapshotDescription () : HBaseProtos.SnapshotDescription 1

org/apache/hadoop/hbase/snapshot/HBaseSnapshotException.getSnapshotDescription:()Lorg/apache/hadoop/hbase/protobuf/generated/HBaseProtos\$SnapshotDescription;

Result	Result
Return value type has been changed from org.apache.hadoop.hbase.protobuf.generated.HBaseProtos.SnapshotDescription to org.apache.hadoop.hbase.client.SnapshotDescription.	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

HBASE-15609: Remove PB references from Result, DoubleColumnInterpreter and any such public facing class for 2.0. hbase-client-1.0.0.jar, Result.class package org.apache.hadoop.hbase.client

[#] Result.getStats () : ClientProtos.RegionLoadStats 1

org/apache/hadoop/hbase/client/Result.getStats:()Lorg/apache/hadoop/hbase/protobuf/generated/ClientProtos\$RegionLoadStats;

Result	Result
Return value type has been changed from org.apache.hadoop.hbase.protobuf.generated.ClientProtos.RegionLoadStats to RegionLoadStats.	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

PrettyPrinter changes

hbase-server-1.0.0.jar, HFilePrettyPrinter.class package org.apache.hadoop.hbase.io.hfile

Result	Result
Return value type has been changed from void to int.	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.