

Cloudera Runtime 7.2.12

Release Notes

Date published: 2021-10-25

Date modified:

CLOUDERA

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

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

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

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

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

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

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

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

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

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

Contents

CVE-2021-45105 & CVE-2021-44832 Remediation for 7.2.12.....	6
CVE-2021-4428 Remediation for 7.2.12.....	6
Overview.....	6
Cloudera Runtime Component Versions.....	6
Using the Cloudera Runtime Maven repository.....	8
Maven Artifacts for Cloudera Runtime 7.2.12.....	8
What's New In Cloudera Runtime 7.2.12.....	28
What's New in Apache Atlas.....	28
What's New in Cruise Control.....	29
What's New in Apache HBase.....	29
What's New in Apache Hive.....	29
What's New in Hue.....	29
What's New in Apache Impala.....	30
What's New in Apache Kudu.....	30
What's New in Schema Registry.....	30
What's New in Cloudera Search.....	31
What's New in Apache Spark.....	31
What's New in Sqoop.....	31
What's New in Streams Replication Manager.....	32
Unaffected Components in this release.....	32
Fixed Issues In Cloudera Runtime 7.2.12.....	32
Fixed Issues in Atlas.....	33
Fixed Issues in Avro.....	33
Fixed Issues in Cloud Connectors.....	33
Fixed issues in Cruise Control.....	33
Fixed issues in Data Analytics Studio.....	33
Fixed Issues in Apache Hadoop.....	33
Fixed Issues in HBase.....	34
Fixed Issues in HDFS.....	34
Fixed Issues in Apache Hive.....	34
Fixed Issues in Hive Warehouse Connector.....	34
Fixed Issues in Hue.....	35
Fixed Issues in Apache Impala.....	35
Fixed Issues in Apache Kafka.....	35
Fixed Issues in Apache Knox.....	35
Fixed Issues in Apache Kudu.....	36
Fixed Issues in Apache Oozie.....	36

Fixed Issues in Phoenix.....	37
Fixed Issues in Parquet.....	37
Fixed Issues in Apache Ranger.....	37
Fixed Issues in Schema Registry.....	38
Fixed Issues in Cloudera Search.....	38
Fixed Issues in Apache Solr.....	38
Fixed Issues in Spark.....	38
Fixed Issues in Apache Sqoop.....	39
Fixed Issues in Streams Messaging Manager.....	39
Fixed Issues in Streams Replication Manager.....	39
Fixed Issues in Apache YARN.....	39
Fixed Issues in Zeppelin.....	40
Fixed Issues in Apache ZooKeeper.....	42
Service Pack in Cloudera Runtime 7.2.12.....	42
Fixed Issues In Cloudera Runtime 7.2.12.7.....	42
Fixed Issues In Cloudera Runtime 7.2.12.8.....	42
Fixed Issues In Cloudera Runtime 7.2.12.9.....	42
Fixed Issues In Cloudera Runtime 7.2.12.10.....	43
Fixed Issues In Cloudera Runtime 7.2.12.11.....	43
Fixed Issues In Cloudera Runtime 7.2.12.12.....	43
Known Issues In Cloudera Runtime 7.2.12.....	43
Known Issues in Apache Atlas.....	43
Known Issues in Apache Avro.....	47
Known issues in Cruise Control.....	47
Known Issues in Data Analytics Studio.....	48
Known Issues in Apache HBase.....	49
Known Issues in HDFS.....	50
Known Issues in Apache Hive.....	51
Known Issues in Hue.....	52
Known Issues in Apache Impala.....	54
Known Issues in Apache Kafka.....	58
Known Issues in Apache Knox.....	60
Known Issues in Apache Kudu.....	60
Known Issues in Apache Oozie.....	60
Known Issues in Apache Phoenix.....	62
Known Issues in Apache Ranger.....	62
Known Issues in Schema Registry.....	63
Known Issues in Cloudera Search.....	64
Known Issues in Apache Spark.....	67
Known Issues for Apache Sqoop.....	68

Known issues in Streams Messaging Manager.....	69
Known Issues in Streams Replication Manager.....	70
Known Issues in MapReduce and YARN.....	73
Known Issues in Apache Zeppelin.....	77
Known Issues in Apache ZooKeeper.....	78
Behavioral Changes In Cloudera Runtime 7.2.12.....	78
Behavioral Changes in Apache Kudu.....	78
Behavioral Changes in Cloudera Search.....	78
Behavioral Changes in Apache Phoenix.....	79
Deprecation Notices In Cloudera Runtime 7.2.12.....	79
Deprecation notices in Apache Kudu.....	80
Deprecation Notices for Apache Kafka.....	80
Deprecation Notices in Apache HBase.....	80

CVE-2021-45105 & CVE-2021-44832 Remediation for 7.2.12

Learn about the CVE-2021-45105 & CVE-2021-44832 Remediation for 7.2.12.

On February 1, 2022, Cloudera released a hotfix to Public Cloud Runtime version 7.2.12. It addresses the CVE and other vulnerability concerns as listed below:

- [CVE-2021-45105](#) which affects Apache Log4j2 versions from 2.0-beta9 to 2.16.0, excluding 2.12.3
- [CVE-2021-44832](#) which affects Apache Log4j2 versions from 2.0-alpha7 to 2.17.0, excluding 2.3.2 and 2.12.4

All new CDP environments with Data Lakes using Runtime 7.2.12 that are registered after this hotfix has been released include the vulnerability fix.

You should upgrade your CDP services running Runtime version 7.2.12 so that they include the hotfix. You can update your existing Data Lake and Data Hubs by performing a maintenance upgrade. You should first upgrade the Data Lake and then upgrade all the Data Hubs that are using the Data Lake. The maintenance upgrade is not supported for RAZ-enabled environments. Refer to [Data Lake upgrade](#) and [Data Hub upgrade](#) documentation.

CVE-2021-4428 Remediation for 7.2.12

You can learn more about the CVE-2021-4428 Remediation for 7.2.12.

On January 3, 2022, Cloudera released Public Cloud runtime version 7.2.12_2. It addresses 2 CVEs and other vulnerability concerns as listed below.

- [CVE-2021-44228](#) which affects Apache Log4j2 versions 2.0 through 2.14.1.
- [CVE-2021-45046](#) which affects Apache Log4j2 version 2.15.0
- [LOGBACK-1591](#) which affects logback versions <= 1.2.7

Cloudera urges all customers on the runtime version 7.2.12 (for Datalake or Datahub) to upgrade their services to the latest version.

Overview

You can review the Release Notes of Cloudera Runtime 7.2.12 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 7.2.12 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	0.11.1.7.2.12.0-291
Apache Atlas	2.1.0.7.2.12.0-291
Apache Calcite	1.21.0.7.2.12.0-291
Apache Avro	1.8.2.7.2.12.0-291

Component	Version
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.2.12.0-291
Apache HBase	2.2.6.7.2.12.0-291
Apache Hive	3.1.3000.7.2.12.0-291
Apache Impala	4.0.0.7.2.12.0-291
Apache Kafka	2.5.0.7.2.12.0-291
Apache Knox	1.3.0.7.2.12.0-291
Apache Kudu	1.15.0.7.2.12.0-291
Apache Livy	0.6.0.7.2.12.0-291
Apache MapReduce	3.1.1.7.2.12.0-291
Apache NiFi	1.13.2.7.2.12.0-291
Apache NiFi Registry	0.8.0.7.2.12.0-291
Apache Oozie	5.1.0.7.2.12.0-291
Apache ORC	1.5.1.7.2.12.0-291
Apache Parquet	1.10.99.7.2.12.0-291
Apache Phoenix	5.1.1.7.2.12.0-291
Apache Ranger	2.1.0.7.2.12.0-291
Apache Solr	8.4.1.7.2.12.0-291
Apache Spark	2.4.8.7.2.12.0-291
Apache Sqoop	1.4.7.7.2.12.0-291
Apache Tez	0.9.1.7.2.12.0-291
Apache Zeppelin	0.8.2.7.2.12.0-291
Apache ZooKeeper	3.5.5.7.2.12.0-291

Other Components

Component	Version
Cruise Control	2.0.100.7.2.12.0-291
Data Analytics Studio	1.4.2.7.2.12.0-291
GCS Connector	2.1.2.7.2.12.0-291
HBase Indexer	1.5.0.7.2.12.0-291
Hue	4.5.0.7.2.12.0-291
Search	1.0.0.7.2.12.0-291
Schema Registry	0.10.0.7.2.12.0-291
Streams Messaging Manager	2.2.0.7.2.12.0-291
Streams Replication Manager	1.1.0.7.2.12.0-291

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.2.12.0-291
Hive Meta Store (HMS)	1.0.0.7.2.12.0-291
Hive on Tez	1.0.0.7.2.12.0-291

Component	Version
Hive Warehouse Connector	1.0.0.7.2.12.0-291
Spark Atlas Connector	0.1.0.7.2.12.0-291
Spark Schema Registry	1.1.0.7.2.12.0-291

Using the Cloudera Runtime Maven repository

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

If you want to build applications or tools for use with Cloudera Runtime components and you are using Maven or Ivy for dependency management, you can pull the Cloudera Runtime artifacts from the Cloudera Maven repository. The repository is available at <https://repository.cloudera.com/artifactory/cloudera-repos/>.



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

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

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

Maven Artifacts for Cloudera Runtime 7.2.12

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	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-classification-updater	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-client-common	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-client-v1	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-client-v2	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-common	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-distro	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-docs	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-graphdb-api	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-graphdb-common	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-intg	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-notification	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-repository	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-server-api	2.1.0.7.2.12.0-291
	org.apache.atlas	atlas-testtools	2.1.0.7.2.12.0-291
	org.apache.atlas	hbase-bridge	2.1.0.7.2.12.0-291
	org.apache.atlas	hbase-bridge-shim	2.1.0.7.2.12.0-291
	org.apache.atlas	hbase-testing-util	2.1.0.7.2.12.0-291
	org.apache.atlas	hdfs-model	2.1.0.7.2.12.0-291
	org.apache.atlas	hive-bridge	2.1.0.7.2.12.0-291
	org.apache.atlas	hive-bridge-shim	2.1.0.7.2.12.0-291
	org.apache.atlas	impala-bridge	2.1.0.7.2.12.0-291
	org.apache.atlas	impala-bridge-shim	2.1.0.7.2.12.0-291
	org.apache.atlas	impala-hook-api	2.1.0.7.2.12.0-291
	org.apache.atlas	kafka-bridge	2.1.0.7.2.12.0-291
	org.apache.atlas	kafka-bridge-shim	2.1.0.7.2.12.0-291
	org.apache.atlas	navigator-to-atlas	2.1.0.7.2.12.0-291
	org.apache.atlas	sample-app	2.1.0.7.2.12.0-291
	org.apache.atlas	sqoop-bridge	2.1.0.7.2.12.0-291
	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.2.12.0-291
Apache Avro	org.apache.avro	avro	1.8.2.7.2.12.0-291
	org.apache.avro	avro-compiler	1.8.2.7.2.12.0-291
	org.apache.avro	avro-ipc	1.8.2.7.2.12.0-291
	org.apache.avro	avro-mapred	1.8.2.7.2.12.0-291
	org.apache.avro	avro-maven-plugin	1.8.2.7.2.12.0-291
	org.apache.avro	avro-protobuf	1.8.2.7.2.12.0-291
	org.apache.avro	avro-service-archetype	1.8.2.7.2.12.0-291
	org.apache.avro	avro-thrift	1.8.2.7.2.12.0-291
	org.apache.avro	avro-tools	1.8.2.7.2.12.0-291
	org.apache.avro	trevni-avro	1.8.2.7.2.12.0-291
	org.apache.avro	trevni-core	1.8.2.7.2.12.0-291
Apache Calcite	org.apache.calcite	calcite-babel	1.21.0.7.2.12.0-291
	org.apache.calcite	calcite-core	1.21.0.7.2.12.0-291
	org.apache.calcite	calcite-druid	1.21.0.7.2.12.0-291
	org.apache.calcite	calcite-kafka	1.21.0.7.2.12.0-291
	org.apache.calcite	calcite-linq4j	1.21.0.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.calcite	calcite-server	1.21.0.7.2.12.0-291
	org.apache.calcite	calcite-ubenchmark	1.21.0.7.2.12.0-291
	org.apache.calcite.avatica	avatica	1.16.0.7.2.12.0-291
	org.apache.calcite.avatica	avatica-core	1.16.0.7.2.12.0-291
	org.apache.calcite.avatica	avatica-metrics	1.16.0.7.2.12.0-291
	org.apache.calcite.avatica	avatica-metrics-dropwizardmetrics	1.16.0.7.2.12.0-291
	org.apache.calcite.avatica	avatica-noop-driver	1.16.0.7.2.12.0-291
	org.apache.calcite.avatica	avatica-server	1.16.0.7.2.12.0-291
	org.apache.calcite.avatica	avatica-standalone-server	1.16.0.7.2.12.0-291
	org.apache.calcite.avatica	avatica-tck	1.16.0.7.2.12.0-291
Apache Druid	org.apache.druid	druid-aws-common	0.17.1.7.2.12.0-291
	org.apache.druid	druid-benchmarks	0.17.1.7.2.12.0-291
	org.apache.druid	druid-console	0.17.1.7.2.12.0-291
	org.apache.druid	druid-core	0.17.1.7.2.12.0-291
	org.apache.druid	druid-gcp-common	0.17.1.7.2.12.0-291
	org.apache.druid	druid-hll	0.17.1.7.2.12.0-291
	org.apache.druid	druid-indexing-hadoop	0.17.1.7.2.12.0-291
	org.apache.druid	druid-indexing-service	0.17.1.7.2.12.0-291
	org.apache.druid	druid-integration-tests	0.17.1.7.2.12.0-291
	org.apache.druid	druid-processing	0.17.1.7.2.12.0-291
	org.apache.druid	druid-server	0.17.1.7.2.12.0-291
	org.apache.druid	druid-services	0.17.1.7.2.12.0-291
	org.apache.druid	druid-sql	0.17.1.7.2.12.0-291
	org.apache.druid	extendedset	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-hive-extensions	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-hbase-security	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-hbase-bloom-filter	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-hbase-datasketches	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-hbase2-extensions	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-hive-google-extensions	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-hdfs-storage	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-histogram	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-kafka-extraction-namespace	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-kafka-indexing-service	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-kerberos	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-kinesis-indexing-service	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-lookups-cached-global	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-lookups-cached-single	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-morc-extensions	0.17.1.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.druid.extensions	druid-parquet-extensions	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-protobuf-extensions	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-h3-extensions	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-stats	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-sql-metadata-storage	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-sql-metadata-storage	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-client-sslcontext	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-contribs-emitter	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-contribs-emitter	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-contribs-extensions	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-contribs-storage	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-cloudfiles-extensions	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-discount	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-dynamodb-extensions	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-dynamodb-emitter	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-dynamodb-sketch	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-dynamodb-average-query	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-dynamodb-emitter	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-distributed-cache	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-distributed-sketch	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-distributed-extensions	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-distributed-in-max	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-distributed-columns	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-distributed-emitter	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-distributed-emitter	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-distributed-view-maintenance	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-distributed-view-selection	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-distributed-metadata-storage	0.17.1.7.2.12.0-291
	org.apache.druid.extensions	druid-distributed-emitter	0.17.1.7.2.12.0-291
GCS Connector	com.google.cloud.bigtable	bigtable-gcs-connector	2.1.2.7.2.12.0-291
	com.google.cloud.bigtable	bigtable-connector	2.1.2.7.2.12.0-291
	com.google.cloud.bigtable	bigtable-gcs	2.1.2.7.2.12.0-291
	com.google.cloud.bigtable	bigtable-aoss	2.1.2.7.2.12.0-291
	com.google.cloud.bigtable	bigtable-distributed	2.1.2.7.2.12.0-291
Apache Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-annotations	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-archives	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-assemblies	3.1.1.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-auth	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-aws	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-azure	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-build-tools	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-client	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-client-api	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-common	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-datajoin	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-distcp	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-extras	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-fs2img	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-gridmix	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-client	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-common	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-config	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-container-service	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-docs	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-hadoop-dependency-client	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-hadoop-dependency-server	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-hadoop-dependency-test	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-interface-admin	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-interface-client	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-interface-server	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-server-framework	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-server-scm	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-test-utils	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdds-tools	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-hdfs	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-kafka	3.1.1.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-kms	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-minicluster	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-minikdc	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-nfs	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-openstack	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-client	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-common	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-csi	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-datanode	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-dist	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-filesystem	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-filesystem-common	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-filesystem-hadoop2	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-filesystem-hadoop3	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-filesystem-shaded	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-insight	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-integration-test	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-interface-client	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-interface-storage	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-network-tests	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-ozon-manager	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-recon	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-reconcodegen	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-s3gateway	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-ozon-tools	1.1.0.7.2.12.0-291
	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-rumen	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-sls	3.1.1.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-streaming	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.2.12.0-291
	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.2.12.0-291
	org.apache.hadoop	mini-chaos-tests	1.1.0.7.2.12.0-291
Apache HBase	org.apache.hbase	hbase-annotations	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-checkstyle	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-client	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-client-project	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-common	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-endpoint	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-examples	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-external-blockcache	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-hadoop-compat	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-hadoop2-compat	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-hbtop	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-http	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-it	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-mapreduce	2.2.6.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.hbase	hbase-metrics	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-metrics-api	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-procedure	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-protocol	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-protocol-shaded	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-replication	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-resource-bundle	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-rest	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-rsgroup	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-server	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-shaded-client	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-shaded-client-project	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-shaded-mapreduce	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-shaded-testing-util	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-shaded-testing-util-tester	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-shell	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-testing-util	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-thrift	2.2.6.7.2.12.0-291
	org.apache.hbase	hbase-zookeeper	2.2.6.7.2.12.0-291
	org.apache.hbase.contrib	hbase-kafka-model	1.0.0.7.2.12.0-291
	org.apache.hbase.contrib	hbase-kafka-proxy	1.0.0.7.2.12.0-291
	org.apache.hbase.contrib	hbase-spark	1.0.0.7.2.12.0-291
	org.apache.hbase.contrib	hbase-spark-it	1.0.0.7.2.12.0-291
	org.apache.hbase.contrib	hbase-spark-protocol	1.0.0.7.2.12.0-291
	org.apache.hbase.contrib	hbase-spark-protocol-shaded	1.0.0.7.2.12.0-291
	org.apache.hbase.contrib	hbase-spark3	1.0.0.7.2.12.0-291
	org.apache.hbase.contrib	hbase-spark3-it	1.0.0.7.2.12.0-291
	org.apache.hbase.contrib	hbase-spark3-protocol	1.0.0.7.2.12.0-291
	org.apache.hbase.contrib	hbase-spark3-protocol-shaded	1.0.0.7.2.12.0-291
	org.apache.hbase.filesystem	hbase-hdfs-impl	1.0.0.7.2.12.0-291
	org.apache.hbase.filesystem	hbase-hdfs	1.0.0.7.2.12.0-291
	org.apache.hbase.thirdparty	hbase-noop-htrace	3.5.0.7.2.12.0-291
	org.apache.hbase.thirdparty	hbase-shaded-gson	3.5.0.7.2.12.0-291
	org.apache.hbase.thirdparty	hbase-shaded-jersey	3.5.0.7.2.12.0-291
	org.apache.hbase.thirdparty	hbase-shaded-jetty	3.5.0.7.2.12.0-291
	org.apache.hbase.thirdparty	hbase-shaded-miscellaneous	3.5.0.7.2.12.0-291
	org.apache.hbase.thirdparty	hbase-shaded-netty	3.5.0.7.2.12.0-291
	org.apache.hbase.thirdparty	hbase-shaded-protobuf	3.5.0.7.2.12.0-291

Project	groupId	artifactId	version
Apache Hive	org.apache.hive	catalogd-unit	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-beeline	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-blobstore	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-classification	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-cli	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-common	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-contrib	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-druid-handler	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-exec	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-hbase-handler	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-hpsql	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-impala	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-it-custom-serde	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-it-druid	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-it-impala	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-it-minikdc	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-it-qfile	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-it-test-serde	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-it-unit	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-it-util	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-jdbc	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-jdbc-handler	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-jmh	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-kryo-registry	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-kudu-handler	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-llap-client	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-llap-common	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-llap-ext-client	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-llap-server	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-llap-tez	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-metastore	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-parser	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-pre-upgrade	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-serde	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-service	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-service-rpc	3.1.3000.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.hive	hive-shims	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-spark-client	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-standalone-metastore	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-storage-api	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-streaming	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-testutils	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-udf	3.1.3000.7.2.12.0-291
	org.apache.hive	hive-vector-code-gen	3.1.3000.7.2.12.0-291
	org.apache.hive	kafka-handler	3.1.3000.7.2.12.0-291
	org.apache.hive.hcatalog	hive-hcatalog-core	3.1.3000.7.2.12.0-291
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	3.1.3000.7.2.12.0-291
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	3.1.3000.7.2.12.0-291
	org.apache.hive.hcatalog	hive-hcatalog-streaming	3.1.3000.7.2.12.0-291
	org.apache.hive.hcatalog	hive-webhcat	3.1.3000.7.2.12.0-291
	org.apache.hive.hcatalog	hive-webhcat-java-client	3.1.3000.7.2.12.0-291
	org.apache.hive.hive-udf-classloader-udf1	hive-udf-classloader-udf1	3.1.3000.7.2.12.0-291
	org.apache.hive.hive-udf-classloader-udf2	hive-udf-classloader-udf2	3.1.3000.7.2.12.0-291
	org.apache.hive.hive-udf-classloader-util	hive-udf-classloader-util	3.1.3000.7.2.12.0-291
	org.apache.hive.hive-udf-vectorized-badexample	hive-udf-vectorized-badexample	3.1.3000.7.2.12.0-291
	org.apache.hive.shims	hive-shims-0.23	3.1.3000.7.2.12.0-291
	org.apache.hive.shims	hive-shims-common	3.1.3000.7.2.12.0-291
	org.apache.hive.shims	hive-shims-scheduler	3.1.3000.7.2.12.0-291
Apache Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.2.12.0-291
Apache Kafka	org.apache.kafka	connect	2.5.0.7.2.12.0-291
	org.apache.kafka	connect-api	2.5.0.7.2.12.0-291
	org.apache.kafka	connect-basic-auth-extension	2.5.0.7.2.12.0-291
	org.apache.kafka	connect-file	2.5.0.7.2.12.0-291
	org.apache.kafka	connect-json	2.5.0.7.2.12.0-291
	org.apache.kafka	connect-mirror	2.5.0.7.2.12.0-291
	org.apache.kafka	connect-mirror-client	2.5.0.7.2.12.0-291
	org.apache.kafka	connect-runtime	2.5.0.7.2.12.0-291
	org.apache.kafka	connect-transforms	2.5.0.7.2.12.0-291
	org.apache.kafka	generator	2.5.0.7.2.12.0-291
	org.apache.kafka	jmh-benchmarks	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-clients	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-cloudera-plugins	2.5.0.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.kafka	kafka-examples	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-log4j-appender	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-examples	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-scala_2.12	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-scala_2.13	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-test-utils	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-upgrade-system-tests-10	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-upgrade-system-tests-11	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-upgrade-system-tests-20	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-upgrade-system-tests-21	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-upgrade-system-tests-22	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-upgrade-system-tests-23	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-streams-upgrade-system-tests-24	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka-tools	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka_2.12	2.5.0.7.2.12.0-291
	org.apache.kafka	kafka_2.13	2.5.0.7.2.12.0-291
Apache Knox	org.apache.knox	gateway-adapter	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-admin-ui	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-applications	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-cloud-bindings	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-demo-ldap	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-discovery-ambari	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-discovery-cm	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-docker	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-i18n	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-performance-test	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-ha	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-jersey	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-rewrite	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-release	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-server	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-server-launcher	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-admin	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-as	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-definitions	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-hbase	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-health	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-hive	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-idbroker	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-impala	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-jkg	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-knoxsso	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-knoxssout	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-livy	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-metadata	1.3.0.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.knox	gateway-service-nifi	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-rm	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-session	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-storm	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-test	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-tgs	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-vault	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-service-webhdfs	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-shell	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-shell-launcher	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-shell-release	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-shell-samples	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-spi	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-test	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-test-idbroker	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-test-release-utils	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-test-utils	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-topology-simple	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-util-common	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-util-configinjector	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-util-launcher	1.3.0.7.2.12.0-291
	org.apache.knox	gateway-util-urltemplate	1.3.0.7.2.12.0-291
	org.apache.knox	hadoop-examples	1.3.0.7.2.12.0-291
	org.apache.knox	knox-cli-launcher	1.3.0.7.2.12.0-291
	org.apache.knox	knox-homepage-ui	1.3.0.7.2.12.0-291
	org.apache.knox	knox-token-management-ui	1.3.0.7.2.12.0-291
	org.apache.knox	webhdfs-kerb-test	1.3.0.7.2.12.0-291
	org.apache.knox	webhdfs-test	1.3.0.7.2.12.0-291
Apache Kudu	org.apache.kudu	kudu-backup-tools	1.15.0.7.2.12.0-291
	org.apache.kudu	kudu-backup2_2.11	1.15.0.7.2.12.0-291
	org.apache.kudu	kudu-backup3_2.12	1.15.0.7.2.12.0-291
	org.apache.kudu	kudu-client	1.15.0.7.2.12.0-291
	org.apache.kudu	kudu-hive	1.15.0.7.2.12.0-291
	org.apache.kudu	kudu-spark2-tools_2.11	1.15.0.7.2.12.0-291
	org.apache.kudu	kudu-spark2_2.11	1.15.0.7.2.12.0-291
	org.apache.kudu	kudu-spark3-tools_2.12	1.15.0.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.kudu	kudu-spark3_2.12	1.15.0.7.2.12.0-291
	org.apache.kudu	kudu-test-utils	1.15.0.7.2.12.0-291
Apache Livy	org.apache.livy	livy-api	0.6.0.7.2.12.0-291
	org.apache.livy	livy-api	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-client-common	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-client-http	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-core_2.11	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-core_2.12	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-examples	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-integration-test	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-repl_2.11	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-repl_2.12	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-rsc	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-scala-api_2.11	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-scala-api_2.12	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-server	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-test-lib	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-thriftserver	0.6.3000.7.2.12.0-291
	org.apache.livy	livy-thriftserver-session	0.6.3000.7.2.12.0-291
Apache Lucene	org.apache.lucene	lucene-analyzers-common	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-analyzers-icu	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-analyzers-kuromoji	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-analyzers-morfologik	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-analyzers-nori	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-analyzers-openslp	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-analyzers-phonetic	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-analyzers-smartcn	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-analyzers-stempel	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-backward-codecs	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-benchmark	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-classification	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-codecs	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-core	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-demo	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-expressions	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-facet	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-grouping	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-highlighter	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-join	8.4.1.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.lucene	lucene-memory	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-misc	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-monitor	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-queries	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-queryparser	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-replicator	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-sandbox	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-spatial	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-spatial-extras	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-spatial3d	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-suggest	8.4.1.7.2.12.0-291
	org.apache.lucene	lucene-test-framework	8.4.1.7.2.12.0-291
Apache Oozie	org.apache.oozie	oozie-client	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-core	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-distro	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-examples	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-server	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-sharelib-git	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-tools	5.1.0.7.2.12.0-291
	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.2.12.0-291
	org.apache.oozie.test	oozie-mini	5.1.0.7.2.12.0-291
Apache ORC	org.apache.orc	orc-core	1.5.1.7.2.12.0-291
	org.apache.orc	orc-examples	1.5.1.7.2.12.0-291
	org.apache.orc	orc-mapreduce	1.5.1.7.2.12.0-291
	org.apache.orc	orc-shims	1.5.1.7.2.12.0-291
	org.apache.orc	orc-tools	1.5.1.7.2.12.0-291
Apache Parquet	org.apache.parquet	parquet-avro	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-cascading	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-cascading3	1.10.99.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.parquet	parquet-column	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-common	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-encoding	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-format-structures	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-generator	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-hadoop	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-jackson	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-pig	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-pig-bundle	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-protobuf	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-scala_2.10	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-thrift	1.10.99.7.2.12.0-291
	org.apache.parquet	parquet-tools	1.10.99.7.2.12.0-291
Apache Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.2	5.1.1.7.2.12.0-291
	org.apache.phoenix	phoenix-client-hbase-2.2	5.1.1.7.2.12.0-291
	org.apache.phoenix	phoenix-connectors-phoenix5-compat	6.0.0.7.2.12.0-291
	org.apache.phoenix	phoenix-core	5.1.1.7.2.12.0-291
	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.1.7.2.12.0-291
	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.1.7.2.12.0-291
	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.1.7.2.12.0-291
	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.1.7.2.12.0-291
	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.1.7.2.12.0-291
	org.apache.phoenix	phoenix-pherf	5.1.1.7.2.12.0-291
	org.apache.phoenix	phoenix-queryserver	6.0.0.7.2.12.0-291
	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.2.12.0-291
	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.2.12.0-291
	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.2.12.0-291
	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.2.12.0-291
	org.apache.phoenix	phoenix-server-hbase-2.2	5.1.1.7.2.12.0-291
	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.2.12.0-291
	org.apache.phoenix	phoenix5-hive	6.0.0.7.2.12.0-291
	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.2.12.0-291
	org.apache.phoenix	phoenix5-spark	6.0.0.7.2.12.0-291
	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.2.12.0-291
	org.apache.phoenix	phoenix5-shaded-commons-cli	1.1.0.7.2.12.0-291
	org.apache.phoenix	phoenix5-shaded-guava	1.1.0.7.2.12.0-291
Apache Ranger	org.apache.ranger	conditions-enrichers	2.1.0.7.2.12.0-291
	org.apache.ranger	credentialbuilder	2.1.0.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.ranger	embeddedwebservice	2.1.0.7.2.12.0-291
	org.apache.ranger	jisql	2.1.0.7.2.12.0-291
	org.apache.ranger	ldapconfigcheck	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-adls-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-atlas-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-atlas-plugin-shim	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-authn	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-distro	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-examples-distro	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-hbase-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-hbase-plugin-shim	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-hdfs-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-hdfs-plugin-shim	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-hive-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-hive-plugin-shim	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-intg	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-kafka-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-kafka-plugin-shim	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-kms	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-kms-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-kms-plugin-shim	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-knox-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-knox-plugin-shim	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-kudu-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-kylin-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-kylin-plugin-shim	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-nifi-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-nifi-registry-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-ozone-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-ozone-plugin-shim	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-plugin-classloader	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-plugins-audit	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-plugins-common	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-plugins-cred	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-plugins-installer	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-raz-adls	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-raz-hook-abfs	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-raz-hook-s3	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-raz-intg	2.1.0.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.ranger	ranger-raz-processor	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-raz-s3	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-raz-s3-lib	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-rms-common	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-rms-hive	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-rms-plugins-common	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-rms-webapp	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-s3-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-sampleapp-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-schema-registry-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-solr-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-solr-plugin-shim	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-sqoop-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-sqoop-plugin-shim	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-storm-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-storm-plugin-shim	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-tagsync	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-tools	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-util	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-yarn-plugin	2.1.0.7.2.12.0-291
	org.apache.ranger	ranger-yarn-plugin-shim	2.1.0.7.2.12.0-291
	org.apache.ranger	sample-client	2.1.0.7.2.12.0-291
	org.apache.ranger	sampleapp	2.1.0.7.2.12.0-291
	org.apache.ranger	shaded-raz-hook-abfs	2.1.0.7.2.12.0-291
	org.apache.ranger	shaded-raz-hook-s3	2.1.0.7.2.12.0-291
	org.apache.ranger	ugsync-util	2.1.0.7.2.12.0-291
	org.apache.ranger	unixauthclient	2.1.0.7.2.12.0-291
	org.apache.ranger	unixauthservice	2.1.0.7.2.12.0-291
	org.apache.ranger	unixusersync	2.1.0.7.2.12.0-291
Apache Solr	org.apache.solr	solr-analysis-extras	8.4.1.7.2.12.0-291
	org.apache.solr	solr-analytics	8.4.1.7.2.12.0-291
	org.apache.solr	solr-cell	8.4.1.7.2.12.0-291
	org.apache.solr	solr-clustering	8.4.1.7.2.12.0-291
	org.apache.solr	solr-core	8.4.1.7.2.12.0-291
	org.apache.solr	solr-dataimporthandler	8.4.1.7.2.12.0-291
	org.apache.solr	solr-dataimporthandler-extras	8.4.1.7.2.12.0-291
	org.apache.solr	solr-jaegertracer-configurator	8.4.1.7.2.12.0-291
	org.apache.solr	solr-langid	8.4.1.7.2.12.0-291
	org.apache.solr	solr-ltr	8.4.1.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.solr	solr-prometheus-exporter	8.4.1.7.2.12.0-291
	org.apache.solr	solr-security-util	8.4.1.7.2.12.0-291
	org.apache.solr	solr-solrj	8.4.1.7.2.12.0-291
	org.apache.solr	solr-test-framework	8.4.1.7.2.12.0-291
	org.apache.solr	solr-velocity	8.4.1.7.2.12.0-291
Apache Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-avro_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-catalyst_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-catalyst_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-core_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-core_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-graphx_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-graphx_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-hadoop-cloud_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-hive_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-hive_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-kubernetes_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-kvstore_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-kvstore_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-launcher_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-launcher_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-mllib-local_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-mllib_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-mllib_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-network-common_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-network-common_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-network-shuffle_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-network-yarn_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-repl_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-repl_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-shaded-raz	3.1.2.7.2.12.0-291
	org.apache.spark	spark-sketch_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-sketch_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.2.12.0-291

Project	groupId	artifactId	version
	org.apache.spark	spark-sql-kafka-0-10_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-sql_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-sql_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-streaming-kafka-0-10_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-streaming_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-streaming_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-tags_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-tags_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-token-provider-kafka-0-10_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-unsafe_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-unsafe_2.12	3.1.2.7.2.12.0-291
	org.apache.spark	spark-yarn_2.11	2.4.8.7.2.12.0-291
	org.apache.spark	spark-yarn_2.12	3.1.2.7.2.12.0-291
Apache Sqoop	org.apache.sqoop	sqoop	1.4.7.7.2.12.0-291
	org.apache.sqoop	sqoop-test	1.4.7.7.2.12.0-291
Apache Tez	org.apache.tez	hadoop-shim	0.9.1.7.2.12.0-291
	org.apache.tez	hadoop-shim-2.8	0.9.1.7.2.12.0-291
	org.apache.tez	tez-api	0.9.1.7.2.12.0-291
	org.apache.tez	tez-aux-services	0.9.1.7.2.12.0-291
	org.apache.tez	tez-common	0.9.1.7.2.12.0-291
	org.apache.tez	tez-dag	0.9.1.7.2.12.0-291
	org.apache.tez	tez-examples	0.9.1.7.2.12.0-291
	org.apache.tez	tez-ext-service-tests	0.9.1.7.2.12.0-291
	org.apache.tez	tez-history-parser	0.9.1.7.2.12.0-291
	org.apache.tez	tez-javadoc-tools	0.9.1.7.2.12.0-291
	org.apache.tez	tez-job-analyzer	0.9.1.7.2.12.0-291
	org.apache.tez	tez-mapreduce	0.9.1.7.2.12.0-291
	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.2.12.0-291
	org.apache.tez	tez-runtime-internals	0.9.1.7.2.12.0-291
	org.apache.tez	tez-runtime-library	0.9.1.7.2.12.0-291
	org.apache.tez	tez-tests	0.9.1.7.2.12.0-291
	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.2.12.0-291
	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.2.12.0-291
	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.2.12.0-291
	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.2.12.0-291

Project	groupId	artifactId	version
Apache Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.2.12.0-291
	org.apache.zeppelin	zeppelin-display	0.8.2.7.2.12.0-291
	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.2.12.0-291
	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.2.12.0-291
	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.2.12.0-291
	org.apache.zeppelin	zeppelin-livy	0.8.2.7.2.12.0-291
	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.2.12.0-291
	org.apache.zeppelin	zeppelin-server	0.8.2.7.2.12.0-291
	org.apache.zeppelin	zeppelin-shaded-raz	0.8.2.7.2.12.0-291
	org.apache.zeppelin	zeppelin-shell	0.8.2.7.2.12.0-291
	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.2.12.0-291
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.5.5.7.2.12.0-291
	org.apache.zookeeper	zookeeper-client-c	3.5.5.7.2.12.0-291
	org.apache.zookeeper	zookeeper-contrib-loggraph	3.5.5.7.2.12.0-291
	org.apache.zookeeper	zookeeper-contrib-rest	3.5.5.7.2.12.0-291
	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.5.5.7.2.12.0-291
	org.apache.zookeeper	zookeeper-docs	3.5.5.7.2.12.0-291
	org.apache.zookeeper	zookeeper-jute	3.5.5.7.2.12.0-291
	org.apache.zookeeper	zookeeper-recipes-election	3.5.5.7.2.12.0-291
	org.apache.zookeeper	zookeeper-recipes-lock	3.5.5.7.2.12.0-291
	org.apache.zookeeper	zookeeper-recipes-queue	3.5.5.7.2.12.0-291

What's New In Cloudera Runtime 7.2.12

You must be aware of the additional functionalities and improvements to features of components in Cloudera Runtime 7.2.12. Learn how the new features and improvements benefit you.

What's New in Apache Atlas

Learn about the new features of Apache Atlas in Cloudera Runtime 7.2.12.

Schema Registry metadata collection

Integrating Schema Registry with Atlas will enable you to persist and view the schemas in Atlas.

See [Schema Registry metadata collection](#) for more information.

Ignore or Prune pattern to filter data

Atlas supports metadata and lineage updates from services like HBase, Hive, Impala, and Spark.

See [Ignore or Prune pattern to filter Hive metadata entities](#) for more information.

What's New in Cruise Control

There are no new features for Cruise Control in Cloudera Runtime 7.2.12.

Cruise Control service is available in CDP Public Cloud

Cruise Control is supported in CDP Public Cloud and Streams Messaging cluster definitions. Beside its basic functionalities, Cruise Control in CDP Public Cloud enables you to scale your Kafka clusters based on your requirements.

What's New in Apache HBase

Learn about the new features of HBase in Cloudera Runtime 7.2.12.

CDPD-28326, 28367: COD - HashTable/SyncTable

HashTable/SyncTable tool is extended to allow alternative token-based authentication mechanism. Data can be synchronized by generating a hash of a table from the source cluster and then only compare the different hashes using the CldrSyncTable command on the target cluster.

What's New in Apache Hive

Learn about the new features of Hive in Cloudera Runtime 7.2.12.

Accessing StorageHandler-based tables

In this release, you provide SERDEPROPERTIES or TBLPROPERTIES when you create the external table. In Ranger > Hadoop SQL and in Hadoop SQL Policies, you set up a StorageHandler policy, and turn on a storage handler authorization property in Cloudera Manager. Hive uses this information to authorize access to the table based on Ranger policies you set up.

For details, see [Accessing StorageHandler and other external tables](#).

Token-based authentication for Cloudera Data Warehouse

Using a token, you can sign on to use Hive and Impala in Cloudera Data Warehouse for a period of time instead of entering your single-sign on (SSO) credentials every time you need to run a query.

For details, see [Token-based authentication for Cloudera Data Warehouse integrations](#).

What's New in Hue

Learn about the new features of Hue in Cloudera Runtime 7.2.12.

Ability to grant fine-grained access to S3 buckets from the S3 File Browser in Hue (Preview)

For better security and ease of use for users, you can create per-user home directories within your S3 bucket and grant fine-grained access to these user directories from the S3 File Browser in Hue. To enable fine-grained access from the Hue S3 File Browser, you must enable Ranger Authorization Service (RAZ) when you register your AWS environment with CDP. For more information, see [Enabling Fine-grained Access Control for Hue S3 File Browser](#).



Note: You need to contact Cloudera to have this feature enabled.

Ability to grant fine-grained access to ADLS Gen2 containers from the ABFS File Browser in Hue (Preview)

For better security and ease of use for users, you can create per-user home directories within your ADLS Gen2 container and grant fine-grained access to these user directories from the ABFS File Browser in Hue. To enable fine-grained access from the Hue ABFS File Browser, you must enable RAZ when you register your Azure environment with CDP. For more information, see [Enabling Fine-grained Access Control from ABFS File Browser in Hue](#).



Note: You need to contact Cloudera to have this feature enabled.

What's New in Apache Impala

There are no new features of Impala in Cloudera Runtime 7.2.12.

What's New in Apache Kudu

Learn about the new features of Kudu in Cloudera Runtime 7.2.12.

Improvements

KUDU-3161: Include FileSystem Path in UUID Mismatch Error

Error logging is enhanced to include filesystem path if the UUID of the instance file in any of the filesystem roots differs from the UUID of other filesystem roots by including the filesystem path. This helps in troubleshooting scenarios where the exact directory is needed.

KUDU-3304: Allow updating table's replication factor via kudu CLI

This patch adds a function to alter table's replication factor, supported in CLI tool, by using the following command:

```
kudu table set_replication_factor  
r [***MASTER_ADDRESSES***] [***TABLE_NAME***] [***REPLICATION_FACTOR***]
```

The CLI tool will return immediately without waiting for the RF change to be effective, new replicas are being up and running when increasing RF, duplicate replicas are being shut down when decreasing RF asynchronously. You can use `kudu cluster ksck` to check whether it has been effective.

What's New in Schema Registry

Learn about the new features of Schema Registry in Cloudera Runtime 7.2.12.

Schema Registry integration with Atlas

You can now enable the integration between Schema Registry and Atlas to view a version of a schema in Atlas.

Atlas serves as a metadata catalog that shows relationships between entities. Additionally, when Kafka is also integrated with Atlas you can connect Kafka topics and schemas.

For more information, see *Integrating with Atlas*.

Related Information

[Integrating with Atlas](#)

What's New in Cloudera Search

Learn about the new features of Cloudera Search in Cloudera Runtime 7.2.12.

The MapReduceIndexerTool now supports ingesting parquet DECIMAL data into Solr

The `ReadAvroParquetFile` morphline command got a new attribute, `decimalConversionEnabled`. When set to true, decimal Parquet data is correctly read (instead of returning raw bytes as earlier). Its default value is false for backwards compatibility.

Decimal data can be stored as a string, float or double Solr field. Only storage as string is lossless, float and double can result in loss of precision.

Optimize phase of MRIT reducers can now be skipped

This release adds a new feature to MapReduceIndexerTool (MRIT) that allows skipping the optimize phase of reducers. To skip the optimize operation at the end of reducers, set the value of the `maxSegments` parameter to 0 or -1.

What's New in Apache Spark

Learn about the new features of Spark in Cloudera Runtime 7.2.12.

Apache Spark 3 version support

- Support for virtual clusters powered by Apache Spark 3 is now available.
- The following functionalities are not currently supported:
 - Deep analysis (visual profiler)
 - HWC - that is, Hive managed ACID tables (Direct Reader & JDBC mode)
 - Phoenix Connector
 - SparkR

See [Running Apache Spark 3 applications](#) and [Data Engineering clusters](#).

What's New in Sqoop

Learn what's new in the Apache Sqoop client in Cloudera Runtime 7.2.12.

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 Replication Manager

Learn about the new features of Streams Replication Manager in Cloudera Runtime 7.2.12.

SRM Service Multi Target Cluster Support

The SRM Service can now target and monitor multiple Kafka clusters. Previously, the Service could only target a single cluster. Targeting multiple clusters enables you to monitor multiple remote Kafka clusters using a single SRM deployment. Configuration is done in Cloudera Manager by adding multiple target Kafka cluster aliases to the Streams Replication Manager Service Target Cluster property. For more information, see [Configuring the service role target cluster](#).

SRM Service Remote Querying [Technical Preview]



Important: Remote Querying is available in this release of Cloudera Runtime but is not ready for production deployment. Cloudera encourages you to explore technical preview features in non-production environments and provide feedback on your experiences through the [Cloudera Community Forums](#).

A new feature, Remote Querying, is introduced for the SRM Service. Remote Querying enables the SRM Service to query other, remote SRM Services and fetch the metrics gathered by the remote SRM Services. This allows users to monitor all replications of a deployment that has multiple instances of SRM with a single SRM Service. For more information, see [Remote Querying](#) and [Configuring Remote Querying](#).

Auto topic creation is disabled by default on source clusters

A new Connector property, `disable.source.topic.auto.creation`, is introduced. This property controls whether auto topic creation is enabled for SRM's internal consumers that fetch data from source clusters. This property is enabled by default meaning that auto topic creation is disabled by default. This property is not directly available for configuration in Cloudera Manager, and does not have a dedicated configuration entry on the UI. Instead, it can be configured through the Streams Replication Manager's Replication Configs Cloudera Manager property. Configuration is possible on a global or replication level.

Unaffected Components in this release

There are no new features for the following components in Cloudera Runtime 7.2.12: Data Analytics Studio – Apache HBase – Apache Hadoop HDFS – Apache Hive – Hue – Apache-Kafka – Apache Knox – Apache Oozie – Apache Phoenix – Apache Ranger – Cloudera Search – Apache Solr – Streams Messaging Manager – Apache ZooKeeper

There are no new features for the following components in Cloudera Runtime 7.2.12:

- Data Analytics Studio
- Apache Hadoop HDFS
- Apache Kafka
- Apache Knox
- Apache Oozie
- Apache Phoenix
- Apache Ranger
- Apache Solr
- Streams Messaging Manager
- Apache ZooKeeper

Fixed Issues In Cloudera Runtime 7.2.12

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.12.

Fixed Issues in Atlas

Review the list of Apache Atlas issues that are resolved in Cloudera Runtime 7.2.12.

CDPD-29335: Random NPE when retrieving tasks.

Handle null pointer exception when retrieving tasks using the admin/tasks endpoint. This issue is now resolved.

Apache patch information

Apache patches in this release. These patches do not have an associated Cloudera bug ID.

- ATLAS-4431
- ATLAS-4398
- ATLAS-4340
- ATLAS-4354
- ATLAS-4348

Fixed Issues in Avro

There are no fixed issues for Avro in Cloudera Runtime 7.2.12.

Fixed Issues in Cloud Connectors

Review the list of Cloud Connectors issues that are resolved in Cloudera Runtime 7.2.12.

CDPD-13652: hdfs copyFromLocal to abfs failed with OutOfMemoryError on DataMart Azure Cluster.

The abfs connector's output streams can buffer blocks to write on disk (default), byte array or off-heap ByteBuffer. Configuration options allow queue length and size of shared pool of uploader threads to set. This issue is now resolved.

Apache patch information

Apache patches in this release. These patches do not have an associated Cloudera bug ID.

- HADOOP-17735
- HADOOP-17195

Fixed issues in Cruise Control

Review the list of Cruise Control issues that are resolved in Cloudera Runtime 7.2.12.

CDPD-23610: Cruise Control script overrides security related properties added to Safety Valve

The issue regarding the Cruise Control script is resolved, the properties added to the Safety Valve are applied for the Cruise Control service.

Fixed issues in Data Analytics Studio

There are no fixed issues for Data Analytics Studio in Cloudera Runtime 7.2.12.

Fixed Issues in Apache Hadoop

There are no fixed issues for Hadoop in Cloudera Runtime 7.2.12.

Apache Patch Information

- HADOOP-17735
- HADOOP-17483
- HADOOP-17304
- HADOOP-17208

Fixed Issues in HBase

There are no fixed issues for HBase in Cloudera Runtime 7.2.12.

Fixed Issues in HDFS

There are no fixed issues for HDFS in Cloudera Runtime 7.2.12.

Apache Patch Information

- None

Fixed Issues in Apache Hive

Review the list of Hive issues that are resolved in Cloudera Runtime 7.2.12.

CDPD-29425: Backport HIVE-25502: Cleaner causes data loss when processing aborted txn with DP.

This issue is now resolved.

CDPD-26114: Backport HIVE-25113: Connection starvation in TxnHandler.isValidWriteIds.

Fixed a connection pool deadlock problem in TxnHandler, which potentially caused connection requests timeout before. This issue is now resolved.

CDPD-25027: FENG: Exception in Calcite while running smart meter query.

Fix cardinality preserving join optimization when a column is backtracked to a constant. This issue is now resolved.

CDPD-20452: Hive - Upgrade jackson to 2.10.5.1 or 2.11.0+ due to CVE-2020-25649.

This issue is now resolved.

Apache Patch Information

- HIVE-25502
- HIVE-25113
- HIVE-24325
- HIVE-24816

Technical Service Bulletins

TSB 2021-532: HWC fails to write empty DataFrame to orc files

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-532: HWC fails to write empty DataFrame to orc files](#)

Fixed Issues in Hive Warehouse Connector

There are no fixed issues for Hive Warehouse Connector in Cloudera Runtime 7.2.12.

Fixed Issues in Hue

There are no fixed issues for Hue in Cloudera Runtime 7.2.12.

Fixed Issues in Apache Impala

There are no fixed issues for Impala in Cloudera Runtime 7.2.12.

Apache Patch Information

- None

Fixed Issues in Apache Kafka

Review the list of Apache Kafka issues that are resolved in Cloudera Runtime 7.2.12.

CDPD-26944: Kafka client entities with null or empty client ID are created in Atlas

Atlas entities with null or empty client ID are no longer created.

CDPD-24428: Topics created with the default replication factor and partition count are incorrectly propagated into Atlas

Topics created with the default replication factor and partition count are now propagated into Atlas with the correct replication factor and partition number.

Fixed Issues in Apache Knox

Review the list of Knox issues that are resolved in Cloudera Runtime 7.2.12.

CDPD-29717: [Logout-knox] From knox homepage clicking logout returns the 500 error code.

Fixed the issue where Knox cookies were not properly deleted after logout. This issue is now resolved.

CDPD-29573: [Logout-knox]On Logout hadoop-jwt cookie should be removed and the description about having previous valid session should be removed.

With this fix now hadoop-jwt cookie is removed after a successful logout. This issue is now resolved.

CDPD-29521: [E2E-HA-Terminate-GCP] 400 Error while getting access token from metadata server with Host does not match SNI.

Make sure the failover request contains proper SNI header. This issue is now resolved.

CDPD-29430: Broken ODBC Connection using HA Datahubs.

Disable HA loadblancing for windows ODBC driver to prevent connection errors. This issue is now resolved.

CDPD-28900: Authentication with IDBroker failed in dex post test runs for DL HA Terminate flow.

Make sure correct SNI header is added after failover. This issue is now resolved.

CDPD-28196: RM UI redirect link to the Spark3 History Server not working.

Spark 3 History Server link Resource Manager UI now works. This issue is now resolved.

Apache patch information

Apache patches in this release. These patches do not have an associated Cloudera bug ID.

- KNOX-2671
- KNOX-2634

Fixed Issues in Apache Kudu

Review the list of Apache Kudu issues that are resolved in Cloudera Runtime 7.2.12.

KUDU-1921: Add ability for clients to require authentication/encryption

Kudu servers support requiring authentication and encryption to be enabled, and clients prefer connecting in a secure way, but if a server does not support authentication and/or encryption, the client will silently connect insecurely, which can lead to a downgrade attack. With this fix, clients can require authentication and encryption to be set using the client API, where if such an attack is attempted, the client will fail to connect to the cluster.

KUDU-2302: Leader crashes if it can't resolve DNS address of a peer

When a tablet replica is elected leader, it constructs Peer objects for each replica in the Raft configuration to send RPCs to each. If during this construction any remote peer cannot be reached for any reason, it could result in a crash. This fix allows to start Peers without a proxy, and retries constructing the proxy the next time a proxy is required.

KUDU-3291: Crash when performing a diff scan after delta flush races with a batch of ops that update the same row

Performing an incremental backup of rows that had many batches of update could result in a crash.

KUDU-3297: KRPC connection negotiation fails with RedHat/CentOS cyrus-sasl-gssapi-2.1.27-5 for secure clusters

Prior to this patch, GSSAPI-involved scenarios of the negotiation-test and security-itest would fail when running against the GSSAPI plugin with patch applied. With this patch, all scenarios in the negotiation-test and the security-itest pass.

Fixed Issues in Apache Oozie

Review the list of Oozie issues that are resolved in Cloudera Runtime 7.2.12.

CDPD-28080: [web UI] Oozie web UI should not serve image from http://extjs.com/s.gif.

OOZIE-3431: Oozie web UI should not serve image from http://extjs.com/s.gif. This issue is now resolved.

CDPD-26335: Oozie server startup error when JDBC URL for a MySQL DB with HA is used.

Oozie did not accept a JDBC url if it contained special characters (e.g.: comma or curly brackets). This is fixed now so Oozie should accept special urls, for example a MySQL HA JDBC url where you typically use comma(s). This issue is resolved.

CDPD-24942: Add a --insecure like parameter to Oozie client so it can ignore certificate errors.

The Oozie client now has a --insecure option which behaves like the curl command's --insecure parameter. With this there's no need to specify the -Djavax.net.ssl.trustStore and -Djavax.net.ssl.trustStorePassword parameters when SSL is enabled. This issue is now resolved.

CDPD-27661: Implement a way to disable the Oozie UI.

OOZIE-3431: There is a new property named "oozie.ui.enabled" users can set for Oozie with a value true or false. By default it's set to true. When set to false, the Oozie UI will be completely disabled, the Oozie server will not even expose the UI resources, hence this can be a workaround for the JQuery vulnerabilities. This issue is now resolved.

CDPD-25864: Support basic authentication in Oozie's callback and notification url.

Oozie's notification mechanism now supports the https://user:password@host:port/... format so it's capable of handling basic authentication. This issue is now resolved.

Apache patch information

- OOZIE-3431

- OOOIE-2136

Fixed Issues in Phoenix

Review the list of Phoenix issues that are resolved in Cloudera Runtime 7.2.12.

OPSAPS-61603: OpsDB upgrade failed.

Fixing minimum memory requirement for OMID TSO service during upgrade to 7.2.12+. This issue is resolved.

Fixed Issues in Parquet

There are no fixed issues for Parquet in Cloudera Runtime 7.2.12.

Apache Patch Information

- None

Fixed Issues in Apache Ranger

Review the list of Ranger issues that are resolved in Cloudera Runtime 7.2.12.

CDPD-28887: Use apt logger to log messages.

Made changes to use the Logger to log the messages by removing System.out.println(...). This issue is resolved.

CDPD-28836: Issue in policies search on report page with user having more than one unix group.

Made changes to fetch all the public group and current user({USER}) assigned policies when user search for policies on report page with username. This issue is resolved.

CDPD-28718: Ranger HiveAuthorizer improvements to handle uncharted hive commands.

This issue is resolved.

CDPD-28535: Ranger Admin: Improve error message while deleting users and groups associated with role.

This patch has changes to improve error message while deleting users and groups which are associated with role(s). This issue is resolved.

CDPD-27464: For ADLS service, User enter password is not secure.

For ADLS service, the password field has type "string" changed this type to "password". This issue is resolved.

CDPD-27323: All policies are exported, when searching reports using roles.

When policies are searched using Role name on the report page. then only specified role name policies are exported. This issue is resolved.

CDPD-20644: Ranger Authorization for StorageHandler based Hive table creation.

This jira provides the functionality of authorizing StorageHandler in the HIVE Create / Alter statement of table. A policy has to be maintained for Storage Type (hbase, phoenix, kafka, jdbc) and corresponding Storage URL for the StorageHandler in Ranger for this Authorization. This issue is resolved.

Apache Patch Information

- RANGER-3468
- RANGER-3454
- RANGER-3285

- RANGER-3463
- RANGER-3372
- RANGER-3368
- RANGER-3361
- RANGER-3336
- RANGER-2950
- RANGER-3374

Fixed Issues in Schema Registry

Review the list of Schema Registry issues that are resolved in Cloudera Runtime 7.2.12.

CDPD-28714: Schema should connect to all Kafka topics with the same name

When having multiple Kafka clusters, with the Schema Registry-Atlas plugin, schemas are connected to all topics with the same name as the schema.

CDPD-26574: SR - Remove the spring security dependency due to CVE-2021-22112, was used by kerberos basic authentication

Removed the dependency on spring-security due to CVE-2021-22112.

CDPD-25052: No active transaction is associated with the thread

Multiple entries of No active transaction in the logs made the logs difficult to read. The logs no longer contain multiple entries.

CDPD-17737: Fix misused exception handling of Schema Registry resources

Endpoint `/api/v1/schemaregistry/schemas/versionsById/{versionId}/branch` responds with 404 Not Found when a `SchemaNotFoundException` is thrown (and not 400 Bad Request).

Fixed Issues in Cloudera Search

There are no fixed issues for Search in Cloudera Runtime 7.2.12.

Fixed Issues in Apache Solr

There are no fixed issues for Solr in Cloudera Runtime 7.2.12.

Apache patch information

- None

Fixed Issues in Spark

Review the list of Spark issues that are resolved in Cloudera Runtime 7.2.12.

CDPD-27961: Fix partition directory double deletes during insert overwrite.

Insert overwrite no longer fails due to partition directory double delete This issue is now resolved.

Apache patch information

Apache patches in this release. These patches do not have an associated Cloudera bug ID.

- None

Fixed Issues in Apache Sqoop

There are no fixed issues for Sqoop in Cloudera Runtime 7.2.12.

Apache patch information

No additional Apache patches.

Fixed Issues in Streams Messaging Manager

Review the list of Streams Messaging Manager issues that are resolved in Cloudera Runtime 7.2.12.

CDPD-27197: Handle requests when KafkaConnect client not defined

When KafkaConnect is not configured, all the associated endpoints (except /is-configured) returns 404.

Fixed Issues in Streams Replication Manager

Review the list of Streams Replication Manager issues that are resolved in Cloudera Runtime 7.2.12.

CDPD-27662: The /remote-topics endpoint reports null source for topics

The SRM Service now properly fills the source field of the topic info provided on the /remote-topics endpoint.

CDPD-14019: SRM may automatically re-create deleted topics on source clusters

Deleted topics are no longer automatically re-created on source clusters.

OPSAPS-60823: Configuring the SRM Client's secure storage is mandatory for unsecured environments

The srm-control tool failed to run in an unsecured cluster. Unnecessary additional configs have been removed from the srm-control tool's configuration when it is used in an unsecured cluster.

OPSAPS-61001: Saving configuration changes for SRM is not possible

The SRM Client's Secure Storage Password property is no longer a mandatory property.

Fixed Issues in Apache YARN

Review the list of YARN issues that are resolved in Cloudera Runtime 7.2.12.

COMPX-6880: YARN-10874 - Refactor NM ContainerLaunch#getEnvDependencies's unit tests.

This issue is now resolved.

COMPX-6716: HDFS-16129 HttpFS signature secret file misuse.

HttpFS used a random secret instead of the configured `hadoop.http.authentication.signature.secret.file` when the deprecated `httpfs.authentication.signature.secret.file` configuration was not set also. This issue is now resolved.

COMPX-6684: YARN-10355 - Refactor NM ContainerLaunch.java#orderEnvByDependencies.

This issue is now resolved.

COMPX-6683: YARN-10814 - YARN should not start with empty `hadoop.http.authentication.signature.secret.file`.

Empty signature secret file - `hadoop.http.authentication.signature.secret.file` - made Yarn's REST endpoints unavailable. From now on Yarn will fall back to random secrets in this case. This issue is now resolved.

COMPX-6627: Add `max-parallel-apps` into YARN scheduler response.

The ResourceManager's Cluster Scheduler API's response was extended with the maximum parallel applications property. This issue is now resolved.

Apache patch information

- YARN-10891

Fixed Issues in Zeppelin

Review the list of Zeppelin issues that are resolved in Cloudera Runtime 7.2.12.

Fixed issues

- CDPD-29838: Zeppelin notebook creation failure due to Hadoop ClassCastException
- CDPD-28834: Credentials file should get saved along with notebook-authorization.json and interpreter.json
- CDPD-24981: Zeppelin notebook can not create table with jdbc phoenix interpreter
- CDPD-23410: Zeppelin notebook_authorized_users::test_only_owners_can_change_permissions test is failing
- CDPD-22469: ZEPPELIN-5231: Livy Interpreter doesn't support Japanese Character - Encoding Issue
- CDPD-17187: Zeppelin - Upgrade to Angular 1.8.0 due to CVEs
- CDPD-20908: Remove log4j-slf4j-impl from JDBC/Hive interpreter
- CDPD-19308: Zeppelin - Upgrade to slf4j 1.7.30
- CDPD-19316: Zeppelin - Upgrade httpclient to 4.5.13+ / 5.0.3+ due to CVE-2020-13956
- CDPD-20461: Zeppelin - Upgrade jackson to 2.10.5.1 or 2.11.0+ due to CVE-2020-25649
- CDPD-20267: Zeppelin build failure on cdpd-master
- CDPD-17471: [ZEPPELIN-5116] Accessing zeppelin via Knox after Knox logout should be redirected to Knox login page
- CDPD-17933: Zeppelin - Upgrade Spring Framework to 4.3.29.RELEASE+ due to CVE-2020-5421
- CDPD-19243: Upgrade to Shiro 1.7.0 due to CVE-2020-17510
- CDPD-18170: Zeppelin - Upgrade or remove auto-value due to shaded Guava CVEs
- CDPD-15497: Harmonize Joda-Time to version 2.10.6 (CDPD harmonized)
- CDPD-17543: Zeppelin UI is not compiling due to corrupted notebooks
- CDPD-16197: Upgrade api-*-1.0.0.jar due to CVEs
- CDPD-16096: Zeppelin - upgrade google-oauth-client to 1.31.0
- CDPD-17017: Upgrade XercesImpl to 2.12.0 due to CVE-2018-2799
- CDPD-16845: Upgrade to Shiro 1.6.0 (CVE-2020-13933)
- CDPD-16111: Upgrade jsoup-1.7.2 (CVE-2015-6748)
- CDPD-16104: Upgrade PostgreSQL JDBC driver to 42.2.16
- CDPD-14569: [ZEPPELIN-4414]. Upgrade Thrift to 0.13
- CDPD-13378: Bump version of Java Native Access (JNA)
- CDPD-16114 Upgrade jackrabbit-webdav 1.5.2 due to CVE-2015-1833
- ZEP-97: [ZEPPELIN-3690] display all columns with the same name in ui-grid
- CDPD-16115: Upgrade JGit due to CVE-2016-5725
- CDPD-14614: Update Spring Framework for Zeppelin (CVE-2018-1270)
- CDPD-11599: Update Quartz Enterprise Job Scheduler for Zeppelin (CVE-2019-13990)
- CDPD-14580: Upgrade Scala for CVEs
- BUG-124121: Password hashing not working in Zeppelin
- CDPD-15628: Compilation failure on dex-box
- CDPD-12920: Upgrade Nimbus-JOSE-JWT to 7.9
- CDPD-14990: Upgrade libpam4j to 1.11 (CVE-2017-12197)
- CDPD-11426: Ensure consistent usage of Jackson to 2.10.3
- CDPD-14579: remove org.reflections (CVE-2020-10683)
- CDPD-14581: Update Spring Framework for Zeppelin in 7.2.1.0 (CVE-2018-1275)
- CDPD-14369: [ZEPPELIN-4736] The use of SslContextFactory is deprecated
- CDPD-11406: Include NOTICE and LICENSE files in component directories

- CDPD-11301: Remove jackson and jersey-bundle
- CDPD-11780: Zeppelin: Remove spark (and other interpreters that are not shipped) source dependencies
- CDPD-11348: Update log4j to address CVE-2019-17571
- CDPD-11501: Update Apache Shiro for Zeppelin to 1.5.3
- CDPD-11571: Zeppelin build failure on cdpd-master due to perfect-scrollbar
- CDPD-10187: Zeppelin - Incorrect version of jackson-mapper-asl in CDP
- CDPD-9119: Zeppelin - Upgrade to Guava 28.1 to avoid CVE-2018-1023
- CDPD-9030: Upgrade jackson-databind to version 2.9.10.3 [CVE-2020-8840]
- CDPD-9454: [ZEPPELIN-4697] Zeppelin scheduler pings terracotta.org
- CDPD-8163: Remove `org.spark-project.hive` dependency
- CDPD-7789: Zeppelin - Upgrade to Jetty 9.4.26 to avoid CVEs
- CDPD-7479: add hadoop-cloud-storage jar in Zeppelin
- CDPD-3600: Sync Zeppelin with community latest version (0.8.2)
- CDPD-2933: [ZEPPELIN-4272] Zeppelin fails to use s3a configured for zeppelin.notebook.dir
- CDPD-1683: KerberosRealm roles should match with local file system, if nothing is specified
- CDPD-2300: Initialize proxyuser with proper configuration
- CDPD-1491: Zeppelin should support doAs
- BUG-120595: [ZEPPELIN-4197] Upgrade Jackson to 2.9.9
- BUG-120606: [ZEPPELIN-4187] Bump up version of Scala from 2.11.8 to 2.11.12 (#3378)
- BUG-120605: [ZEPPELIN-4186] Bump up version of org.jsoup:jsoup (#3377)
- BUG-120596: [ZEPPELIN-4188] Upgrade Jetty to 9.4.18.v20190429
- BUG-120594: ZEPPELIN-4193 Upgrade Bouncy Castle bcpkix-jdk15on to 1.60
- BUG-120593: [ZEPPELIN-4185] Upgrade Thrift to 0.12.0 (#3376)
- CDPD-1009: ZEPPELIN-4168: Use secure URLs for Maven repositories (#3370)
- CDPD-717: [Zeppelin 3792] Zeppelin SPNEGO support
- ZEP-79: Disable fs.file.impl cache to ensure RawLocalFS is used
- BUG-109581: [ZEPPELIN-3741] Do not clear "Authorization" header if Z-server is running behind proxy
- BUG-106906: Add shiro-tools-hasher in Zeppelin
- BUG-106297: JDBC interpreter log file is missing in zeppelin log directory
- BUG-102172: Include Google Connector in Zeppelin
- BUG-98604: Correct tutorial link should be added in interpreter page
- BUG-100845: Remove livy2.pyspark3 interpreter on zeppelin side
- BUG-114354: Fixes to make s3 storage work
- BUG-114354: Change Zeppelin to use unshaded jars
- BUG-103954: Exclude other dependencies in STS shaded JDBC driver to prevent conflict
- BUG-103715: fix handshake_failure download
- CDPD-10288: Zeppelin Notebook Initialisation fails with CNF error in RAZ Enabled Cluster

Apache patch information

Apache patches in this release. These patches do not have an associated Cloudera bug ID.

- ZEPPELIN-5231
- ZEPPELIN-5116
- ZEPPELIN-3741
- ZEPPELIN-4168
- ZEPPELIN-4185
- ZEPPELIN-4193
- ZEPPELIN-4188
- ZEPPELIN-4186
- ZEPPELIN-4187

- ZEPPELIN-4197
- ZEPPELIN-4272
- ZEPPELIN-4697
- ZEPPELIN-4736
- ZEPPELIN-3690
- ZEPPELIN-4414

Fixed Issues in Apache ZooKeeper

Review the list of ZooKeeper issues that are resolved in Cloudera Runtime 7.2.12.

OPSAPS-61738: ZooKeeper data directory (/var/lib/zookeeper) no longer has world readable permission

This issue is now resolved.

Service Pack in Cloudera Runtime 7.2.12

You can review the list of CDP Public Cloud hotfixes rolled into Cloudera Runtime 7.2.12. This will help you to verify if a hotfix provided to you on a previous CDP Public Cloud release was included in this release.

- HOTFIX-5069

Fixed Issues In Cloudera Runtime 7.2.12.7

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.12.7.

The following issues are resolved:

- HOTREQ-950 IDBroker client excessively adds SSL client config causing OOM issues
- HOTREQ-964 Release: HIVE-25574: Replace clob with varchar when storing creation metadata

Fixed Issues In Cloudera Runtime 7.2.12.8

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.12.8.

The following issues are resolved:

- HOTREQ-1036 Bug Fix for SPARK-39083
- HOTREQ-1091 ITAU Casting invalid dates does not produce NULL
- HOTREQ-1114 Hue does not work with medium duty DL because IDBroker config has comma separated URLs

Fixed Issues In Cloudera Runtime 7.2.12.9

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.12.9.

The following issues are resolved:

- HOTREQ-1161 - CFM - NIFI nodes disconnect frequently
- HOTREQ-1178 - CDPD-44832 - HUE Oozie workflow rerun fails.

Fixed Issues In Cloudera Runtime 7.2.12.10

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.12.10.

The following issues are resolved:

- HOTREQ-1224 - include CSA-3742 into CSA-DH
- HOTREQ-1201 - HADOOP-18476 - ABFS and S3A FileContext bindings to close wrapped filesystems in finalizer
- HOTREQ-1222 - ABFS - Disable readAhead for 7.2.12 and higher versions

Fixed Issues In Cloudera Runtime 7.2.12.11

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.12.11.

CVE

- Upgrade Apache Commons Text to 1.10.0 due to CVE-2022-42889

Fixed Issues In Cloudera Runtime 7.2.12.12

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.12.12.

The following issues are resolved:

- HOTREQ-1320 - HOTFIX for Bug - Add delegation token support for long running spark job
- HOTREQ-1321 - Hotfix for CVE-2022-25168 in CDP 7.2.11 version

Known Issue CDPD-52789

CDPD-52789: After performing the Datalake upgrade, HMS service goes to error state. The cluster event logs show the following error: Cloudera Manager reported health issues with node(s):...[The following services are in bad health: hive...]

Using Cloudera Manager, restart the Hive MetaStore service. Wait for few minutes, the error clears automatically.

Known Issues In Cloudera Runtime 7.2.12

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.2.12.

Known Issues in Apache Atlas

Learn about the known issues in Apache Atlas, the impact or changes to the functionality, and the workaround.
DOCS-12597: NiFi service user's Atlas interaction (entity_read operations) filling Ranger audit collectionNone.

Problem: NiFi frequently interacts with the Atlas API, which generates Audit logs of the API calls. These Audit logs are indexed in SOLR, which is backed by statically sized EBS volumes on the SDX core nodes (3*250 GB). The amount of these logs causes the EBS volumes to fill up, which takes the SOLR service down

Workaround: To reduce disk space consumption:

1. Reduce the SOLR TTL of the Ranger-audit collection
2. Increase the EBS volumes to their maximum of 1TB
3. Apply a cm_atlas audit filter to filter READ events from the NiFi service user against the Atlas API.

DOCS-12697: Atlas canary check is disabled by default

Problem: Check for Atlas was disabled by default, hence Cloudera Manager might not show proper health alerts.

Workaround: You must manually enable the canary check from Atlas configurations

CDPD-24089: Atlas creates HDFS path entities for GCP path and the qualified name of those entities does not have a cluster name appended.

None

CDPD-22082: ADLS Gen2 metadata extraction: If the queue is not cleared before performing Incremental extraction, messages are lost.

After successfully running Bulk extraction, you must clear the queue before running Incremental extraction.

CDPD-19996: Atlas AWS S3 metadata extractor fails when High Availability is configured for IDBroker.

If you have HA configured for IDBroker, make sure your cluster has only one IDBroker address in core-site.xml. If your cluster has two IDBroker addresses in core-site.xml, remove one of them, and the extractor must be able to retrieve the token from IDBroker.

CDPD-19798: Atlas /v2/search/basic API does not retrieve results when the search text mentioned in the entity filter criteria (like searching by Database or table name) has special characters like + - & ! () { } [] ^ " ~ * ? :

You can invoke the API and mention the search string (with special characters) in the query attribute in the search parameters.

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 "at" character (@), Atlas does not return the expected results or generate appropriate search suggestions.

Consider leaving out the portion of the search string that includes the @ sign, 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

CDPD-24058: The Atlas-Kafka hook creates a new entity instead of linking them

When the import-kafka.sh tool is used and later the plugin is enabled in Kafka configurations, new incomplete topic entities is created. The tool is not linking the existing topics with the clients.

None

CDPD-29663: Error while connecting topic with schema in Atlas

The error occurred when Schema Registry tried to make a relationship in Atlas between a schema and a non-existent corresponding topic.

Known Issues in Apache Avro

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

CDPD-23451: Remove/replace jackson-mapper-asl dependency.

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

None.

Known issues in Cruise Control

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

CDPD-47616: Unable to initiate rebalance, number of valid windows (NumValidWindows) is zero

If a Cruise Control rebalance is initiated with the rebalance_disk parameter and Cruise Control is configured to fetch metrics from Cloudera Manager, Cruise Control stops collecting metrics from the partitions that are moved. This is because Cloudera Manager does not collect metrics from moved partitions due to an issue in Kafka (KAFKA-10320).

If the metrics are not available, the partition is considered invalid by Cruise Control. This results in Cruise Control blocking rebalance operations and proposal generation.

Rolling restart Kafka broker roles.

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 run 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.

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 `postOperation` is implemented only for `postDeleteColumn()`.
- If Create Table fails, full rights for that table persist for the user who attempted to create it. If another user later succeeds in creating the table, the user who made the failed attempt still has the full rights.

Workaround: N/A

Apache Issue: [HBASE-6992](#)

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 2023-667: HBase snapshot export failure can lead to data loss

When using Replication Manager for Apache HBase (HBase) snapshot replication, data loss will occur if both of the following conditions are met: (i) the external account used for the operation has delete access to the target storage location, and (ii) the snapshot export fails. If these conditions are

met, the cleanup operation, which is automatically performed after the failure, would delete all data in the root folder of the snapshot, not only the snapshot files. If the user account does not have the delete permission on the target folder, the data remains unaffected.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2023-667: HBase snapshot export failure can lead to data loss](#)

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](#))

Technical Service Bulletins

TSB 2021-531: HDFS Namenode checkpoints not running after installing or upgrading to CDP 7.1.7 in a TLS enabled cluster

HDFS Namenode checkpoints will not run correctly after installing or upgrading to CDP 7.1.7 in a TLS enabled cluster. Recent CDP Public Cloud releases are also affected, see full list of affected CDP releases below.

In an HDFS cluster with High Availability, the Standby NameNode (SBN) is responsible for periodically generating a new checkpoint (also called an FsImage). The standby NameNode then uploads this checkpoint to the Active NameNode (ANN). The purpose of the checkpoint is to reduce the NameNode's startup time by reducing the number of edit logs that must be replayed on NameNode restart.

The image upload from the SBN to ANN is performed over the HTTP protocol using an embedded Jetty HTTP server in the ANN. Due to a bug in the Jetty version 9.4.39 shipped with CDP, this checkpoint upload can fail and the ANN will not receive new FsImage files. The following error is seen in the SBN service logs.

```
java.io.IOException: Exception during image upload
```

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-531: HDFS Namenode checkpoints not running after installing or upgrading to CDP 7.1.7 in a TLS enabled cluster](#)

TSB 2023-666: Out of order HDFS snapshot deletion may delete renamed/moved files, which may result in data loss

Cloudera has discovered a bug in the Apache Hadoop Distributed File System (HDFS) snapshot implementation. Deleting an HDFS snapshot may incorrectly remove files in the .Trash directories or remove renamed files from the current file system state. This is an unexpected behavior because deleting an HDFS snapshot should only delete the files stored in the specified snapshot, but not data in the current state.

In the particular HDFS installation in which the bug was discovered, deleting one of the snapshots caused certain files to be moved to trash and deletion of some of the files in a .Trash directory. Although it is clear that the conditions of the bug are (1) out-of-order snapshot deletion and (2) files moved to trash or other directories, we were unable to replicate the bug in other HDFS installations after executing similar test operations with a variety of different sequences. We also did not observe any actual data loss in our tests. However, there is a remote possibility that this bug may lead to data loss.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2023-666: Out of order HDFS snapshot deletion may delete renamed/moved files, which may result in data loss](#)

Known Issues in Apache Hive

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

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 `valuepath-to-ssl-client.xml`.

Technical Service Bulletins

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.

Unable to delete, move, or rename directories within the S3 bucket from Hue

You may not be able to rename, move, or delete directories within your S3 bucket from the Hue web interface. This is because of an underlying issue, which will be fixed in a future release.

You can move, rename, or delete a directory using the HDFS commands as follows:

1. SSH into your CDP environment host.
2. To delete a directory within your S3 bucket, run the following command:

```
hdfs dfs -rm -r [***COMPLETE-PATH-TO-S3-BUCKET**]/[***DIREC
TORY-NAME***]
```

3. To rename a folder, create a new directory and run the following command to move files from the source directory to the target directory:

```
hdfs dfs -mkdir [***DIRECTORY-NAME***]
```

```
hdfs dfs -mv [***COMPLETE-PATH-TO-S3-BUCKET**]/[***SOURCE-D
IRECTORY***] [***COMPLETE-PATH-TO-S3-BUCKET**]/[***TARGET-D
IRECTORY***]
```

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

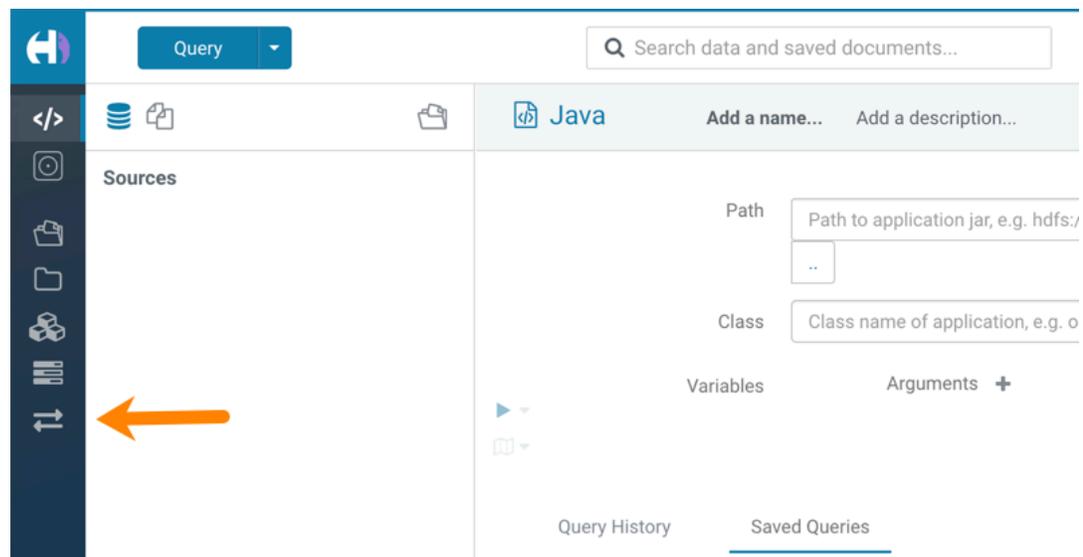
Hue cannot handle columns containing non-UTF8 data. As a result, you may see the following error while queying tables from the Impala editor in Hue: 'utf8' codec can't decode byte 0x91 in position 6: invalid start byte.

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-xxxxxxxx-xxxxxxxxxxxx-xxxxxxxx
.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.

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

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

None.

Known Issues in Apache Impala

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

Impala known limitation when querying compacted tables

When the compaction process deletes the files for a table from the underlying HDFS location, the Impala service does not detect the changes as the compactions does not allocate new write ids.

When the same table is queried from Impala it throws a 'File does not exist' exception that looks something like this:

```
Query Status: Disk I/O error on <node>:22000: Failed to open HDFS
file hdfs://nameservice1/warehouse/tablespace/managed/hive/<database>/<table>/xxxxx
Error(2): No such file or directory Root cause: RemoteException:
File does not exist: /warehouse/tablespace/managed/hive/<database>/<table>/xxxx
```

Use the [REFRESH/INVALIDATE](#) statements on the affected table to overcome the 'File does not exist' exception.

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.

IMPALA-2603: Crash: impala::Coordinator::ValidateCollectionSlots

A query could encounter a serious error if includes multiple nested levels of INNER JOIN clauses involving subqueries.

IMPALA-3094: Incorrect result due to constant evaluation in query with outer join

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 2021-479: 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 Apache 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`.

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

CDPD-29307: Kafka producer entity stays in incomplete state in Atlas

Atlas creates incomplete Kafka client entities that are postfixed with the metadata namespace.

None

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 Hosts Configuration.
 - b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

Enter the following to turn off the collection of partition level metrics:

```
[KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_entity_update_enabled=false
```

Replace [KAFKA_SERVICE_NAME] with the service name of Kafka obtained in step 1. The service name should always be in lower case.

- c. Click Save Changes.

Known Issues in Apache Knox

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

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.

CDPD-29302: The Atlas lineage information is missing in case of HWC JDBC write.

None

CDPD-29297: HWC + Oozie issue: Cannot create PoolableConnectionFactory

Currently only Spark cluster mode is supported in the Oozie Spark Action with Hive Warehouse Connector (HWC).

Use Spark action in cluster mode.

```
<spark xmlns="uri:oozie:spark-action:1.0">
  ...
  <mode>cluster</mode>
  ...
</spark>
```

CDPD-26975: Using the ABFS / S3A connectors in an Oozie workflow where the operations are "secured" may trigger an IllegalArgumentException with the error message java.net.URISyntaxException: Relative path in absolute URI.

Set the following XML configuration in the Datahub cluster's Cloudera Manager:

1. In the Cloudera Manager Admin Console, go to the Oozie service.
2. Click the Configuration tab.
3. In the Oozie Server Advanced Configuration Snippet (Safety Valve) for oozie-site.xml field, set the following:

Set the following if you are using Amazon S3:

```
<property>
  <name>oozie.service.HadoopAccessorService.fs.s3a</name>
  <value>fs.s3a.buffer.dir=/tmp/s3a</value>
</property>
```

Set the following if you are using ABFS:

```
<property>
  <name>oozie.service.HadoopAccessorService.fs.abfs</name>
  <value>fs.azure.buffer.dir=/tmp/abfs</value>
</property>

<property>
  <name>oozie.service.HadoopAccessorService.fs.abfss</name>
  <value>fs.azure.buffer.dir=/tmp/abfss</value>
</property>
```

4. Enter a Reason for change, and then click Save Change to commit the changes.
5. Restart the Oozie service.

Oozie 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 7.2.12.

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.

CDPD-17962: Ranger roles do not work when you upgrade from any CDP Private Cloud Base to CDP Private cloud base. Roles which are created prior to upgrade work as expected, issue is only for new roles created post upgrade and authorization enforced via ranger policies wont work for these new roles. This behavior is only observed with the upgraded cluster; a newly installed cluster does not show this behavior.

There are two possible workarounds to resolve this issue:

1. Update database entries (Recommended):
 - `select * from x_ranger_global_state where state_name='RangerRole';`
 - `update x_ranger_global_state set app_data='{ "Version": "2" }' where state_name='RangerRole';`

Or

2. Add a property in safety valve under ranger-admin-site which will bypass the `getAppDataVersion` method:

Known Issues in Schema Registry

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

CDPD-29663: Error while connecting topic with schema in Atlas

When Atlas integration is enabled for Schema Registry but there is no Kafka topic with the same name as the schema then an error will be logged in the schema registry logs. The status of the record in the `atlas_events` table will also be set to failed. This can be ignored because the status was actually successful. The issue will be fixed in CDP 7.2.13.

CDPD-29465: SR Confluent API's /compatibility endpoint returns invalid payload

Confluent-compatible API /compatibility returns an invalid payload. The field `compatible` cannot be parsed, it must be `is_compatible`. The issue will be fixed in CDP 7.2.13.

CDPD-49217 and CDPD-50309: Schema Registry caches user group membership indefinitely

Schema Registry caches the Kerberos user and group information indefinitely and does not catch up on group membership changes.

Restart Schema Registry after group membership changes.

CDPD-56890: New schemas cannot be created following an upgrade

If you delete the latest version of a schema (the one with the highest ID) from the Schema Registry database before an upgrade, you might not be able to create new schemas after you upgrade the cluster to a newer version.



Important: In CDP Public Cloud, this issue only manifests when upgrading from Cloudera Runtime 7.2.12 or lower to 7.2.14 or higher.

1. Access the Schema Registry database. Go to Cloudera Manager Schema Registry Configuration and search for "database" if you don't know the name, host, or port of the Schema Registry database.
2. Cross reference the ID's in the `schemaVersionId` column of the `schmema_version_state` table with the ID's found in the `schema_version_info` table.
3. Delete all records from the `schema_version_state` table that contains a `schemaVersionId` not present in the `schema_version_info` table.

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

Schema Registry is unable to create entities in Atlas if Schema Registry uses an Oracle database. The following will be present in the Schema Registry log if you are affected by this issue:

```
ERROR com.cloudera.dim.atlas.events.AtlasEventsProcessor: An error occurred while processing Atlas events.
java.lang.IllegalArgumentException: Cannot invoke com.hortonworks.registries.schemaregistry.AtlasEventStorable.setType on bean class 'class com.hortonworks.registries.schemaregistry.AtlasEventStorable' - argument type mismatch - had objects of type "java.lang.Long" but expected signature "java.lang.Integer"
```

This issue causes the loss of audit data on Oracle environments.

None.

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

HBase Lily indexer REST port does not support SSL

When using the `--http` argument for the `hbase-indexer` command line tool to invoke Lily indexer through REST API, you can add/list/remove indexers with any user without the need for authentication.

Switch off the REST API setting the `hbaseindexer.httpserver.disabled` environment parameter to true (by default this is false). This switches off the REST interface, so no one can use the `--http` argument when using the `hbase-indexer` command line tool. This also means that users need to authenticate as `hbase` user in order to use the `hbase-indexer` tool.

The Solr admin UI is only accessible with full `solr_admin` permission in Ranger

Full `solr_admin` permission is required in Ranger to access the Solr Admin UI.

None.

Indexing fails with `socketTimeout`

Starting from CDH 6.0, the HTTP client library used by Solr has a default socket timeout of 10 minutes. Because of this, if a single request sent from an indexer executor to Solr takes more than 10 minutes to be serviced, the indexing process fails with a timeout error.

This timeout has been raised to 24 hours. Nevertheless, there still may be use cases where even this extended timeout period proves insufficient.

If your `MapreduceIndexerTool` or `HBaseMapreduceIndexerTool` batch indexing jobs fail with a timeout error during the `go-live` (Live merge, `MERGEINDEXES`) phase (This means the merge takes longer than 24 hours).

Use the `--go-live-timeout` option where the timeout can be specified in milliseconds.

If the timeout occurs during Near real time (NRT) indexing, Cloudera suggests you try the following workarounds:

- Check the batch size of your indexing job. Sending too large batches to Solr might increase the time needed on the Solr server to process the incoming batch.
- If your indexing job uses `deleteByQuery` requests, consider using `deleteById` wherever possible as `deleteByQuery` involves a complex locking mechanism on the Solr side which makes processing the requests slower.
- Check the number of executors for your Spark Crunch Indexer job. Too many executors can overload the Solr service. You can configure the number of executors by using the `--mappers` parameter

- Check that your Solr installation is correctly sized to accommodate the indexing load, making sure that the number of Solr servers and the number of shards in your target collection are adequate.
- The socket timeout for the connection can be configured in the morphline file. Add the `solrClientSocketTimeout` parameter to the `solrLocator` command

Example

```
SOLR_LOCATOR :
{
  collection : test_collection
  zkHost : "zookeeper1.example.corp:2181/solr"
  # 10 minutes in milliseconds
  solrClientSocketTimeout: 600000
  # Max number of documents to pass per RPC from morphline to
  Solr Server
  # batchSize : 10000
}
```

Splitshard operation on HDFS index checks local filesystem and fails

When performing a shard split on an index that is stored on HDFS, `SplitShardCmd` still evaluates free disk space on the local file system of the server where Solr is installed. This may cause the command to fail, perceiving that there is no adequate disk space to perform the shard split.

Run the following command to skip the check for sufficient disk space altogether:

- On nonsecure clusters:

```
curl 'http://[***SOLR_SERVER_HOSTNAME***]:8983/solr/admin/collections?action=SPLITSHARD&collection=[***COLLECTION_NAME***]&shard=[***SHARD_TO_SPLIT***]&skipFreeSpaceCheck=true'
```

- On secure clusters:

```
curl -k -u : --negotiate 'http://[***SOLR_SERVER_HOSTNAME***]:8985/solr/admin/collections?action=SPLITSHARD&collection=[***COLLECTION_NAME***]&shard=[***SHARD_TO_SPLIT***]&skipFreeSpaceCheck=true'
```

Replace `[***SOLR_SERVER_HOSTNAME***]` with a valid Solr server hostname, `[***COLLECTION_NAME***]` with the collection name, and `[***SHARD_TO_SPLIT***]` with the ID of the to split.

To verify that the command executed successfully, check overseer logs for a similar entry:

```
2021-02-02 12:43:23.743 INFO (OverseerThreadFactory-9-thread-5-processing-n:myhost.example.com:8983_solr) [c:example s:shard1] o.a.s.c.a.c.SplitShardCmd Skipping check for sufficient disk space
```

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.

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.

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 collection --create`.

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

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`.

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.

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.

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

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

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.

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).

Known Issues in Apache Spark

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

CDPD-30637: Spark Error: ClassNotFoundException: org.apache.hadoop.hive.llap.io.api.LlapProxy

Execute the following command: `sudo ln -s /opt/cloudera/parcels/CDH/jars/hive-llap-client-3.1.3000.7.2.12.0-291.jar /opt/cloudera/parcels/CDH/lib/spark3/jars/`. The link to the jar must be created on all gateway and nodemanager nodes.

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: Me
taException(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/tablesp
ace/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 for Streams Messaging Manager in Cloudera Runtime 7.2.12.

CDPD-28002: The Cluster Replications tab is missing from the SMM UI

By default the Cluster Replications tab is visible in the SMM UI when the SMM and SRM services are integrated with each other. However, due to an issue with one of the SMM Knox rewrite rules, the Cluster Replications tab is not rendered in the SMM UI. As a result of this, the Cluster Replications tab cannot be accessed in the UI and replication related metrics cannot be viewed.

Manually update the rewrite.xml file.

1. SSH into the DataHub host where the Knox Gateway Service is located.
2. Go to `/opt/cloudera/parcels/CDH/lib/knox/data/services/smm-ui/2.1.0/`.

The rewrite.xml file is located in this directory.

3. Edit the rewrite.xml file.

You must add two entries. These are as follows:

- Nest the following `<rule>` element under the `<rules>` element.

```
<rule dir="OUT" name="SMM-UI/smm-ui/outbound/apiV2">
  <rewrite template="{ $frontend[path] }/smm-ui/api/v2/admin/" />
</rule>
```

- Nest the following `<apply>` element under the `<content type="application/javascript">` element.

```
<apply path="/api/v2/admin/" rule="SMM-UI/smm-ui/outbound/apiV2" />
```

4. Restart Knox.

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

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

SMM cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You would have to override bootstrap server URL (hostname:port as set in the listeners for broker). Add the bootstrap server details in SMM safety valve in the following path:

Cloudera Manager SMM Configuration Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml Add the following value for bootstrap servers Save Changes Restart SMM :

```
streams.messaging.manager.kafka.bootstrap.servers=<comma-separated list of brokers>
```

OPSAPS-59597: SMM UI logs are not supported by Cloudera Manager

Cloudera Manager does not support the log type used by SMM UI.

View the SMM UI logs on the host.

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:...

Additional security properties must be set for SMM.

- In Cloudera Manager, select the SMM service.
- Go to Configuration.
- 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 Streams Replication Manager

Learn about the known issues in Streams Replication Manager, the impact or changes to the functionality, and the workaround.

Known Issues

CDPD-22089: SRM does not sync re-created source topics until the offsets have caught up with target topic

Messages written to topics that were deleted and re-created are not replicated until the source topic reaches the same offset as the target topic. For example, if at the time of deletion and re-creation there are 100 messages on the source and target clusters, new messages will only get replicated once the re-created source topic has 100 messages. This leads to messages being lost.

None

CDPD-11079: Blacklisted topics appear in the list of replicated topics

If a topic was originally replicated but was later disallowed (blacklisted), it will still appear as a replicated topic under the /remote-topics REST API endpoint. As a result, if a call is made to this endpoint, the disallowed topic will be included in the response. Additionally, the disallowed topic will also be visible in the SMM UI. However, its Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

CDPD-30275: SRM may automatically re-create deleted topics on target clusters

If auto.create.topics.enable is enabled, deleted topics might get automatically re-created on target clusters. This is a timing issue. It only occurs if remote topics are deleted while the replication of the topic is still ongoing.

- Remove the topic from the topic allowlist with srm-control. For example:

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1]
```

- Wait until SRM is no longer replicating the topic.
- Delete the remote topic in the target cluster.

OPSAPS-61814: Using the service dependency method to configure Kerberos enabled co-located clusters is not supported

Using the Streams Replication Manager Co-located Kafka Cluster Alias property to auto-configure the connection to a Kerberos enabled co-located Kafka cluster is not supported. In a case like this, the generated JAAS configuration contains host-specific configuration. This causes SRM to fail to connect to the co-located Kafka cluster on other hosts.

Define your co-located Kafka clusters using Kafka credentials. For more information, see [Defining co-located Kafka clusters using Kafka credentials](#). Alternatively, use the Streams Replication Manager's Replication Configs property to configure the connection to the co-located Kafka clusters.

OPSAPS-63992: Rolling restart unavailable for SRM

Initiating a rolling restart for the SRM service is not possible. Consequently, performing a rolling upgrade of the SRM service is also not possible.

None

CDPD-31745: SRM Control fails to configure internal topic when target is earlier than Kafka 2.3

When the target Kafka cluster of a replication is earlier than version 2.3, the srm-control internal topic is created with an incorrect configuration (cleanup.policy=compact). This causes the srm-control topic to lose the replication filter records, causing issues in the replication.

After a replication is enabled where the target Kafka cluster is earlier than 2.3, manually configure all srm-control.[***SOURCE CLUSTER ALIAS***].internal topics in the target cluster to use cleanup.policy=compact.

CDPD-31235: Negative consumer group lag when replicating groups through SRM

SRM checkpointing reads the offset-syncs topic to create offset mappings for committed consumer group offsets. In some corner cases, it is possible that a mapping is not available in offset-syncs. In a case like this SRM simply copies the source offset, which might not be a valid offset in the replica topic.

One possible situation is if there is an empty topic in the source cluster with a non-zero end offset (for example, retention already removed the records), and a consumer group which has a committed offset set to the end offset. If replication is configured to start replicating this topic, it will not have an offset mapping available in offset-syncs (as the topic is empty), causing SRM to copy the source offset.

This can cause issues when automatic offset synchronization is enabled, as the consumer group offset can be potentially set to a high number. SRM never rewinds these offsets, so even when there is a correct offset mapping available, the offset will not be updated correctly.

After offset mappings are created, stop the consumers of the group and set the committed offsets of the group to the end of the topic on the target cluster with this command:

```
kafka-consumer-groups --bootstrap-server [***HOST***]:[***PORT***] --group [***GROUP***] --topic [***SOURCE CLUSTER ALIAS***].[***TOPIC***] --reset-offsets --to-latest --execute
```

Alternatively, set it to the beginning of the topic with this command:

```
kafka-consumer-groups --bootstrap-server [***HOST***]:[***PORT***] --group <group> --topic [***SOURCE CLUSTER ALIAS***].[***TOPIC***] --reset-offsets --to-earliest --execute
```

OPSAPS-67772: SRM Service metrics processing fails when the noexec option is enabled for /tmp

The SRM Service role uses /tmp to extract RocksDB .so files, which are required for metrics processing to function. If the noexec option is enabled for the /tmp directory, the SRM Service role is not able load the required RocksDB files. This results in metrics processing failing.

1. In Cloudera Manager, select the SRM service and go to Configuration.

2. Add the following to SRM Service Environment Advanced Configuration Snippet (Safety Valve). Do this for all SRM Service role instances.

```
ROCKSDB_SHAREDLIB_DIR=[ ***PATH*** ]
```

Replace `[***PATH***]` with a directory that is not `/tmp`.

OPSAPS-67738: SRM Service role's Remote Querying feature does not work when the noexec option is enabled for /tmp

The SRM Service role puts the Netty native libraries into the `/tmp` directory. As a result, If the `noexec` option is enabled for the `/tmp` directory, the Remote Querying feature will fail to function.

1. In Cloudera Manager, select the SRM service and go to Configuration.
2. Add the following to `SRM_JVM_PERF_OPTS`.

```
-Dio.netty.native.workdir=[ ***PATH*** ]
```

Replace `[***PATH***]` with a directory that is not `/tmp`.

OPSAPS-62546: Kafka External Account SSL keypassword configuration is used incorrectly by SRM

When a Kafka External Account specifies a keystore that uses an SSL key password, SRM uses it as the `ssl.keystore.key` configuration. Due to using the incorrect `ssl.keystore.key` configuration, SRM will fail to load the keystore in certain cases.

Workaround: For the keystores used by the Kafka External Accounts, the SSL key password should match the SSL keystore password, and the SSL keystore key password should not be provided. Alternatively, you can use the legacy connection configurations based on the `streams.replication.manager.configs` to specify the SSL key password.

Limitations

SRM cannot replicate Ranger authorization policies to or from Kafka clusters

Due to a limitation in the Kafka-Ranger plugin, SRM cannot replicate Ranger policies to or from clusters that are configured to use Ranger for authorization. If you are using SRM to replicate data to or from a cluster that uses Ranger, disable authorization policy synchronization in SRM. This can be achieved by clearing the Sync Topic Acls Enabled (`sync.topic.acls.enabled`) checkbox.

SRM cannot ensure the exactly-once semantics of transactional source topics

SRM data replication uses at-least-once guarantees, and as a result cannot ensure the exactly-once semantics (EOS) of transactional topics in the backup/target cluster.



Note: Even though EOS is not guaranteed, you can still replicate the data of a transactional source, but you must set `isolation.level` to `read_committed` for SRM's internal consumers. This can be done by adding `[***SOURCE CLUSTER ALIAS***]->[***TARGET CLUSTER ALIAS***].consumer.isolation.level=read_committed` to the Streams Replication Manager's Replication Configs SRM service property in Cloudera Manger.

SRM checkpointing is not supported for transactional source topics

SRM does not correctly translate checkpoints (committed consumer group offsets) for transactional topics. Checkpointing assumes that the offset mapping function is always increasing, but with transactional source topics this is violated. Transactional topics have control messages in them, which take up an offset in the log, but they are never returned on the consumer API. This causes the mappings to decrease, causing issues in the checkpointing feature. As a result of this limitation, consumer failover operations for transactional topics is not possible.

Known Issues in MapReduce and YARN

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

Known Issues

COMPX-5817: Queue Manager UI will not be able to present a view of pre-upgrade queue structure. CM Store is not supported and therefore Yarn will not have any of the pre-upgrade queue structure preserved.

When a Data Hub cluster is deleted, all saved configurations are also deleted. All YARN configurations are saved in CM Store and this is yet to be supported in Data Hub and Cloudera Manager. Hence, the YARN queue structure also will be lost when a Data Hub cluster is deleted or upgraded or restored.

COMPX-5244: Root queue should not be enabled for auto-queue creation

After dynamic auto child creation is enabled for a queue using the YARN Queue Manager UI, you cannot disable it using the YARN Queue Manager UI. That can cause problem when you want to switch between resource allocation modes, for example from weight mode to relative mode. The YARN Queue Manager UI does not let you to switch resource allocation mode if there is at least one dynamic auto child creation enabled parent queue in your queue hierarchy.

If the dynamic auto child creation enabled parent queue is NOT the root or the root.default queue: Stop and remove the dynamic auto child creation enabled parent queue. Note that this stops and remove all of its child queues as well.

If the dynamic auto child creation enabled parent queue is the root or the root.default queue: You cannot stop and remove neither the root nor the root.default queue. You have to change the configuration in the applicable configuration file:

1. In Cloudera Manager, navigate to YARN>>Configuration.
2. Search for capacity scheduler and find the Capacity Scheduler Configuration Advanced Configuration Snippet (Safety Valve) property.
3. Add the following configuration: `yarn.scheduler.capacity.<queue-path>.auto-queue-creation-v2.enabled=false` For example: `yarn.scheduler.capacity.root.default.auto-queue-creation-v2.enabled=false` Alternatively, you can remove the `yarn.scheduler.capacity.<queue-path>.auto-queue-creation-v2.enabled` property from the configuration file.
4. Restart the Resource Manager.

COMPX-5589: Unable to add new queue to leaf queue with partition capacity in Weight/Absolute mode Scenario

1. User creates one or more partitions.
2. Assigns a partition to a parent with children
3. Switches to the partition to distribute the capacities
4. Creates a new child queue under one of the leaf queues but the following error is displayed:

```
Error :
2021-03-05 17:21:26,734 ERROR
com.cloudera.cpx.server.api.repositories.SchedulerRepository: Val
idation failed for Add queue
operation. Error message: CapacityScheduler configuration vali
dation failed:java.io.IOException:
Failed to re-init queues : Parent queue 'root.test2' have childr
en queue used mixed of weight
mode, percentage and absolute mode, it is not allowed, please do
uble check, details:
{Queue=root.test2.test2childNew, label= uses weight mode}. {Que
ue=root.test2.test2childNew,
```

```
label=partition uses percentage mode}
```

To create new queues under leaf queues without hitting this error, perform the following:

1. Switch to Relative mode
 2. Create the required queues
 3. Create the required partitions
 4. Assign partitions and set capacities
 5. Switch back to Weight mode
1. Create the entire queue structure
 2. Create the required partitions
 3. Assign partition to queues
 4. Set partition capacities

COMPX-5264: Unable to switch to Weight mode on creating a managed parent queue in Relative mode

In the current implementation, if there is an existing managed queue in Relative mode, then conversion to Weight mode is not be allowed.

To proceed with the conversion from Relative mode to Weight mode, there should not be any managed queues. You must first delete the managed queues before conversion. In Weight mode, a parent queue can be converted into managed parent queue.

COMPX-5549: Queue Manager UI sets maximum-capacity to null when you switch mode with multiple partitions

If you associate a partition with one or more queues and then switch the allocation mode before assigning capacities to the queues, an Operation Failed error is displayed as the `max-capacity` is set to null.

After you associate a partition with one or more queues, in the YARN Queue Manager UI, click Overview > <Partition name> from the dropdown list and distribute capacity to the queues before switching allocation mode or creating placement rules.

COMPX-4992: Unable to switch to absolute mode after deleting a partition using YARN Queue Manager

If you delete a partition (node label) which has been associated with queues and those queues have capacities configured for that partition (node label), the CS.xml still contains the partition (node label) information. Hence, you cannot switch to absolute mode after deleting the partition (node label).

It is recommended not to delete a partition (node label) which has been associated with queues and those queues have capacities configured for that partition (node label).

COMPX-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 ResourceManagers

When YARN High Availability is enabled there are multiple ResourceManagers. Queue Manager receives multiple ResourceManager URLs for a High Availability cluster. It picks the active ResourceManager URL only when Queue Manager page is loaded. Queue Manager cannot handle it gracefully when the currently active ResourceManager goes down while the user is still using the Queue Manager UI.

Reload the Queue Manager page manually.

COMPX-3329: Autorestart is not enabled for Queue Manager in Data Hub

In a Data Hub cluster, Queue Manager is installed with autorestart disabled. Hence, if Queue Manager goes down, it will not restart automatically.

If Queue Manager goes down in a Data Hub cluster, you must go to the Cloudera Manager Dashboard and restart the Queue Manager service.

Third party applications do not launch if MapReduce framework path is not included in the client configuration

MapReduce application framework is loaded from HDFS instead of being present on the NodeManagers. By default the `mapreduce.application.framework.path` property is set to the appropriate value, but third party applications with their own configurations will not launch.

Set the `mapreduce.application.framework.path` property to the appropriate configuration for third party applications.

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 include the environment variables HADOOP_HOME, PATH, LANG, and TZ in the allowlist, but the container launch environments do not have these variables set up automatically.

You can manually add the required environment variables to the allowlist 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.

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:cdrwa 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:

- Application Timeline Server (ATSV2 and ATSV1)
- Container Resizing
- Distributed or Centralized Allocation of Opportunistic Containers
- Distributed Scheduling
- Docker on YARN (DockerContainerExecutor) on Data Hub clusters
- Fair Scheduler
- GPU support for Docker
- Hadoop Pipes
- Native Services
- Pluggable Scheduler Configuration
- Queue Priority Support
- Reservation REST APIs
- Resource Estimator Service
- Resource Profiles
- (non-Zookeeper) ResourceManager State Store
- Rolling Log Aggregation
- Shared Cache
- YARN Federation
- Moving jobs between queues

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 7.2.12

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 7.2.12.

Behavioral Changes in Apache Kudu

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

Summary:

The logic to select a particular time source when the pseudo-source 'auto' used for --time_source is changed.

Previous behavior:

The 'system' time source would be auto-selected if a Kudu server with --time_source=auto is run in an environment where the instance detector is not aware of a dedicated NTP server. Those are usually available through the host-only interface, at least for EC2 and GCE instances.

New behavior:

The 'builtin' time source would be auto-selected if a Kudu server runs with --time_source=auto in an environment where the instance detector isn't aware of dedicated NTP servers and the --builtin_ntp_servers flag is set to a valid value. Otherwise, if --builtin_ntp_servers flag is set to an empty/invalid value, 'system' becomes the auto-selected time source for platforms supporting get_ntptime() API, otherwise the catch-all case is used to auto-select 'system_unsync' as the time source.

Behavioral Changes in Cloudera Search

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

Summary:

Accessing Solr Admin UI in CDP clusters secured with Ranger

Previous behavior:

The Solr Admin UI was accessible for users with non-admin privileges.

New behavior

Only users with admin privileges can access the Solr Admin UI.

Summary:

Invalid Atomic Update operations now fail

Previous behavior:

Invalid Atomic Updates threw a warning message.

New behavior:

Invalid Atomic Updates fail with an Exception.

Summary:

Admin API address has changed

Previous behavior:

In Solr 7 both `curl -k --negotiate -u: "https://hostname -f :8985/solr/?op=GETDELEGATION TOKEN"` and `curl -k --negotiate -u: "https://hostname -f :8985/solr/admin?op=GETDELEGATION TOKEN"` commands worked.

New behavior

In Solr 8 only `curl -k --negotiate -u: "https://hostname -f :8985/solr/admin?op=GETDELEGATION TOKEN"` command (with the 'admin' string added) works.

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.

Deprecation Notices In Cloudera Runtime 7.2.12

Certain features and functionalities have been removed or deprecated in Cloudera Runtime 7.2.12. 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 7.2.12. 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 Apache Kafka are deprecated or removed in Cloudera Runtime 7.2.12. 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 7.2.12. 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>void postCloneSnapshot (ObserverContext<MasterCoprocesorEnvironment>, HBaseProtos.SnapshotDescription, HTableDescriptor)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>void postCreateTable (ObserverContext<MasterCoprocesorEnvironment>, HTableDescriptor, HRegionInfo[])</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>void postDeleteSnapshot (ObserverContext<MasterCoprocesorEnvironment>, HBaseProtos.SnapshotDescription)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>void postGetTableDescriptors (ObserverContext<MasterCoprocesorEnvironment>, List<HTableDescriptor>)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>void postModifyTable (ObserverContext<MasterCoprocesorEnvironment>, TableName, HTableDescriptor)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>void postRestoreSnapshot (ObserverContext<MasterCoprocesorEnvironment>, HBaseProtos.SnapshotDescription, HTableDescriptor)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>void postSnapshot (ObserverContext<MasterCoprocesorEnvironment>, HBaseProtos.SnapshotDescription, HTableDescriptor)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>void preCloneSnapshot (ObserverContext<MasterCoprocesorEnvironment>, HBaseProtos.SnapshotDescription, HTableDescriptor)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>void preCreateTable (ObserverContext<MasterCoprocesorEnvironment>, HTableDescriptor, HRegionInfo[])</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>void preDeleteSnapshot (ObserverContext<MasterCoprocesorEnvironment>, HBaseProtos.SnapshotDescription)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>void preGetTableDescriptors (ObserverContext<MasterCoprocesorEnvironment>, List<TableName>, List<HTableDescriptor>)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>void preModifyTable (ObserverContext<MasterCoprocesorEnvironment>, TableName, HTableDescriptor)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>void preRestoreSnapshot (ObserverContext<MasterCoprocesorEnvironment>, HBaseProtos.SnapshotDescription, HTableDescriptor)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>void preSnapshot (ObserverContext<MasterCoprocesorEnvironment>, HBaseProtos.SnapshotDescription, HTableDescriptor)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.

RegionObserver interface changes

The following changes are introduced to the `RegionObserver` interface.

[#] interface `RegionObserver` (13)

Change	Result
Abstract method <code>void postCloseRegionOperation (ObserverContext<RegionCoprocesorEnvironment>, HRegion.Operation)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>void postCompactSelection (ObserverContext<RegionCoprocesorEnvironment>, Store, ImmutableList<StoreFile>)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.

Abstract method <code>voidpostCompactSelection (ObserverContext<RegionCoprocesorEnvironment>, Store, ImmutableList<StoreFile>, CompactionRequest)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>voidpostGetClosestRowBefore (ObserverContext<RegionCoprocesorEnvironment>, byte[], byte[], Result)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>DeleteTrackerpostInstantiateDeleteTracker (ObserverContext<RegionCoprocesorEnvironment>, DeleteTracker)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>voidpostSplit (ObserverContext<RegionCoprocesorEnvironment>, HRegion, HRegion)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>voidpostStartRegionOperation (ObserverContext<RegionCoprocesorEnvironment>, HRegion.Operation)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>StoreFile.ReaderpostStoreFileReaderOpen (ObserverContext<RegionCoprocesorEnvironment>, FileSystem, Path, FSDataInputStreamWrapper, long, CacheConfig, Reference, StoreFile.Reader)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>voidpostWALRestore (ObserverContext<RegionCoprocesorEnvironment>, HRegionInfo, HLogKey, WALEdit)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>InternalScannerpreFlushScannerOpen (ObserverContext<RegionCoprocesorEnvironment>, Store, KeyValueScanner, InternalScanner)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>voidpreGetClosestRowBefore (ObserverContext<RegionCoprocesorEnvironment>, byte[], byte[], Result)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>StoreFile.ReaderpreStoreFileReaderOpen (ObserverContext<RegionCoprocesorEnvironment>, FileSystem, Path, FSDataInputStreamWrapper, long, CacheConfig, Reference, StoreFile.Reader)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>voidpreWALRestore (ObserverContext<RegionCoprocesorEnvironment>, HRegionInfo, HLogKey, WALEdit)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.

WALObserver interface changes

The following changes are introduced to the WALObserver interface:

[#] interface WALObserver

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

Server API changes

[#] class RpcServer (12)

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

Following abstract methods are removed:

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

Replication and WAL changes

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

[#] classWALKey (8)

Change	Result
Access level of field clusterIds has been changed from protected to private.	A client program may be interrupted by IllegalAccessException exception.
Access level of field compressionContext has been changed from protected to private.	A client program may be interrupted by IllegalAccessException exception.
Access level of field encodedRegionName has been changed from protected to private.	A client program may be interrupted by IllegalAccessException exception.
Access level of field tablename has been changed from protected to private.	A client program may be interrupted by IllegalAccessException 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 <code>createTableAsync (HTableDescriptor, byte[] [])</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>disableTableAsync (TableName)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>enableTableAsync (TableName)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>getCompactionState (TableName)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>getCompactionStateForRegion (byte[])</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>isSnapshotFinished (HBaseProtos.SnapshotDescription)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>snapshot (String, TableName, HBaseProtos.SnapshotDescription.Type)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>snapshot (HBaseProtos.SnapshotDescription)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>takeSnapshotAsync (HBaseProtos.SnapshotDescription)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> 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 <code>void</code> to <code>java.util.concurrent.Future<java.lang.Void></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.

[#] `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 <code>void</code> to <code>java.util.concurrent.Future<java.lang.Void></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.

`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 <code>void</code> to <code>java.util.concurrent.Future<java.lang.Void></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.

`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 <code>org.apache.hadoop.hbase.protobuf.generated.AdminProtos.GetRegionInfoResponse.CompactionState</code> to <code>CompactionState</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.

[#] `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 org.apache.hadoop.io.WritableComparable<HColumnDescriptor>.	A client program may be interrupted by NoSuchMethodError exception.

class HTableDescriptor (3)

Change	Result
Removed super-interface org.apache.hadoop.io.WritableComparable<HTableDescriptor>.	A client program may be interrupted by NoSuchMethodError exception.
Field META_TABLEDESC of type HTableDescriptor has been removed from this class.	A client program may be interrupted by NoSuchFieldError 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 HColumnDescriptor[] to client.ColumnFamilyDescriptor[].	This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

[#] interface Table (4)

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

class 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 getNextKeyHint (KeyValue) has been removed from this class.	A client program may be interrupted by NoSuchMethodError 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 <code>IncompatibleClassChangeError</code> or <code>InstantiationException</code> 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 <code>java.util.Map<java.lang.String,master.RegionState></code> to <code>java.util.List<master.RegionState></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.

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 <code>org.apache.hadoop.hbase.protobuf.generated.HBaseProtos.SnapshotDescription</code> to <code>org.apache.hadoop.hbase.client.SnapshotDescription</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.

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 <code>org.apache.hadoop.hbase.protobuf.generated.ClientProtos.RegionLoadStats</code> to <code>RegionLoadStats</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.

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.