

Cloudera Runtime 7.2.9

Release Notes

Date published: 2021-04-30

Date modified:

CLOUdera

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

Legal Notice

© Cloudera Inc. 2025. All rights reserved.

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

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

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

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

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

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

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

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

Contents

CVE-2021-45105 & CVE-2021-44832 Remediation for 7.2.9.....	5
CVE-2021-4428 Remediation for 7.2.9.....	5
Overview.....	5
Cloudera Runtime Component Versions.....	5
Using the Cloudera Runtime Maven repository.....	7
Maven Artifacts for Cloudera Runtime 7.2.9.0.....	7
What's New In Cloudera Runtime 7.2.9.....	27
What's New in Apache Hive.....	27
What's New in Hue.....	27
What's New in Apache Impala.....	27
What's New in Apache Kafka.....	27
What's New in Apache Kudu.....	28
What's New in Apache Phoenix.....	29
What's New in Schema Registry.....	29
What's New in Apache Spark.....	30
What's New in Sqoop.....	30
What's new in Streams Messaging Manager.....	30
What's New in Apache Hadoop YARN.....	31
Unaffected Components in this release.....	32
Fixed Issues In Cloudera Runtime 7.2.9.....	32
Fixed Issues in Atlas.....	33
Fixed Issues in Cloud Connectors.....	36
Fixed issues in Cruise Control.....	36
Fixed issues in Data Analytics Studio.....	36
Fixed Issues in Apache Hadoop.....	36
Fixed Issues in HBase.....	37
Fixed Issues in HDFS.....	38
Fixed Issues in Apache Hive.....	38
Fixed Issues in Hue.....	39
Fixed Issues in Apache Impala.....	39
Fixed Issues in Apache Kafka.....	40
Fixed Issues in Apache Knox.....	40
Fixed Issues in Apache Kudu.....	43
Fixed Issues in Apache Oozie.....	44
Fixed Issues in Phoenix.....	47
Fixed Issues in Parquet.....	47
Fixed Issues in Apache Ranger.....	47

Fixed Issues in Schema Registry.....	50
Fixed Issues in Cloudera Search.....	50
Fixed Issues in Spark.....	50
Fixed Issues in Apache Sqoop.....	52
Fixed Issues in Streams Messaging Manager.....	53
Fixed Issues in Streams Replication Manager.....	53
Fixed Issues in Apache YARN.....	53
Fixed Issues in Zeppelin.....	54
Fixed Issues in Apache ZooKeeper.....	54
 Fixed Issues In Cloudera Runtime 7.2.9.12.....	 54
 Service Pack in Cloudera Runtime 7.2.9.....	 55
 Known Issues In Cloudera Runtime 7.2.9.....	 55
Known Issues in Apache Atlas.....	55
Known Issues in Apache Avro.....	58
Known Issues in Data Analytics Studio.....	58
Known Issues in Apache HBase.....	59
Known Issues in HDFS.....	61
Known Issues in Apache Hive.....	61
Known Issues in Hue.....	63
Known Issues in Apache Impala.....	65
Known Issues in Apache Kafka.....	69
Known Issues in Kerberos.....	71
Known Issues in Apache Knox.....	71
Known Issues in Apache Kudu.....	71
Known Issues in Apache Oozie.....	72
Known Issues in Apache Phoenix.....	72
Known Issues in Apache Ranger.....	72
Known Issues in Schema Registry.....	73
Known Issues in Cloudera Search.....	73
Known Issues in Apache Solr.....	77
Known Issues in Apache Spark.....	77
Known issues in Streams Messaging Manager.....	78
Known Issues in Streams Replication Manager.....	79
Known Issues for Apache Sqoop.....	82
Known Issues in MapReduce and YARN.....	82
Known Issues in Apache Zeppelin.....	88
Known Issues in Apache ZooKeeper.....	88
 Behavioral Changes In Cloudera Runtime 7.2.9.....	 88
Behavioral Changes in Apache Phoenix.....	88
 Deprecation Notices In Cloudera Runtime 7.2.9.....	 89
Deprecation notices in Apache Kudu.....	89
Deprecation Notices for Apache Kafka.....	90
Deprecation Notices in Apache HBase.....	90

CVE-2021-45105 & CVE-2021-44832 Remediation for 7.2.9

Learn about the CVE-2021-45105 & CVE-2021-44832 Remediation for 7.2.9.

On February 3, 2022, Cloudera released a hotfix to Public Cloud Runtime version 7.2.9. 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 7.2.9 that are registered after this hotfix has been released include the vulnerability fix.

You should upgrade your CDP services running Runtime version 7.2.9 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 7.2.9

You can learn more about the CVE-2021-4428 Remediation for 7.2.9.

On January 3, 2022, Cloudera released Public Cloud runtime version 7.2.9_6. 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 7.2.9 (for Datalake or Datahub) to upgrade their services to the latest version.

Overview

You can review the Release Notes of Cloudera Runtime 7.2.9 for release-specific information related to new features and improvements, bug fixes, deprecated features and components, known issues, and changed features that can affect product behavior.

Cloudera Runtime Component Versions

You must be familiar with the versions of all the components in the Cloudera Runtime 7.2.9 distribution to ensure compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.8.0.7.2.9.0-203
Apache Atlas	2.1.0.7.2.9.0-203
Apache Calcite	1.21.0.7.2.9.0-203
Apache Avro	1.8.2.7.2.9.0-203

Component	Version
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.2.9.0-203
Apache HBase	2.2.6.7.2.9.0-203
Apache Hive	3.1.3000.7.2.9.0-203
Apache Impala	4.0.0.7.2.9.0-203
Apache Kafka	2.5.0.7.2.9.0-203
Apache Knox	1.3.0.7.2.9.0-203
Apache Kudu	1.14.0.7.2.9.0-203
Apache Livy	0.6.0.7.2.9.0-203
Apache MapReduce	3.1.1.7.2.9.0-203
Apache NiFi	1.13.2.7.2.9.0-203
Apache NiFi Registry	0.8.0.7.2.9.0-203
Apache Oozie	5.1.0.7.2.9.0-203
Apache ORC	1.5.1.7.2.9.0-203
Apache Parquet	1.10.99.7.2.9.0-203
Apache Phoenix	5.1.1.7.2.9.0-203
Apache Ranger	2.1.0.7.2.9.0-203
Apache Solr	8.4.1.7.2.9.0-203
Apache Spark	2.4.7.7.2.9.0-203
Apache Sqoop	1.4.7.7.2.9.0-203
Apache Tez	0.9.1.7.2.9.0-203
Apache Zeppelin	0.8.2.7.2.9.0-203
Apache ZooKeeper	3.5.5.7.2.9.0-203

Other Components

Component	Version
Data Analytics Studio	1.4.2.7.2.9.0-203
GCS Connector	1.9.10.7.2.9.0-203
HBase Indexer	1.5.0.7.2.9.0-203
Hue	4.5.0.7.2.9.0-203
Search	1.0.0.7.2.9.0-203
Schema Registry	0.10.0.7.2.9.0-203
Streams Messaging Manager	2.1.0.7.2.9.0-203
Streams Replication Manager	1.0.0.7.2.9.0-203

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.2.9.0-203
Hive Meta Store (HMS)	1.0.0.7.2.9.0-203
Hive on Tez	1.0.0.7.2.9.0-203
Hive Warehouse Connector	1.0.0.7.2.9.0-203

Component	Version
Spark Atlas Connector	0.1.0.7.2.9.0-203
Spark Schema Registry	1.1.0.7.2.9.0-203

Using the Cloudera Runtime Maven repository

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

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



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

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

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

Maven Artifacts for Cloudera Runtime 7.2.9.0

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

Project	groupId	artifactId	version
Apache Accumulo	org.apache.accumulo	accumulo-core	1.7.0.7.2.9.0-203
	org.apache.accumulo	accumulo-examples-simple	1.7.0.7.2.9.0-203
	org.apache.accumulo	accumulo-fate	1.7.0.7.2.9.0-203
	org.apache.accumulo	accumulo-gc	1.7.0.7.2.9.0-203
	org.apache.accumulo	accumulo-master	1.7.0.7.2.9.0-203
	org.apache.accumulo	accumulo-minicluster	1.7.0.7.2.9.0-203
	org.apache.accumulo	accumulo-monitor	1.7.0.7.2.9.0-203
	org.apache.accumulo	accumulo-proxy	1.7.0.7.2.9.0-203
	org.apache.accumulo	accumulo-server-base	1.7.0.7.2.9.0-203
	org.apache.accumulo	accumulo-shell	1.7.0.7.2.9.0-203
	org.apache.accumulo	accumulo-start	1.7.0.7.2.9.0-203
	org.apache.accumulo	accumulo-test	1.7.0.7.2.9.0-203
	org.apache.accumulo	accumulo-trace	1.7.0.7.2.9.0-203
	org.apache.accumulo	accumulo-tracer	1.7.0.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.accumulo	accumulo-tserver	1.7.0.7.2.9.0-203
Apache Atlas	org.apache.atlas	atlas-authorization	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-classification-updater	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-client-common	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-client-v1	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-client-v2	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-common	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-distro	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-docs	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-graphdb-api	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-graphdb-common	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-intg	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-notification	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-repository	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-server-api	2.1.0.7.2.9.0-203
	org.apache.atlas	atlas-testtools	2.1.0.7.2.9.0-203
	org.apache.atlas	hbase-bridge	2.1.0.7.2.9.0-203
	org.apache.atlas	hbase-bridge-shim	2.1.0.7.2.9.0-203
	org.apache.atlas	hbase-testing-util	2.1.0.7.2.9.0-203
	org.apache.atlas	hdfs-model	2.1.0.7.2.9.0-203
	org.apache.atlas	hive-bridge	2.1.0.7.2.9.0-203
	org.apache.atlas	hive-bridge-shim	2.1.0.7.2.9.0-203
	org.apache.atlas	impala-bridge	2.1.0.7.2.9.0-203
	org.apache.atlas	impala-bridge-shim	2.1.0.7.2.9.0-203
	org.apache.atlas	impala-hook-api	2.1.0.7.2.9.0-203
	org.apache.atlas	kafka-bridge	2.1.0.7.2.9.0-203
	org.apache.atlas	kafka-bridge-shim	2.1.0.7.2.9.0-203
	org.apache.atlas	navigator-to-atlas	2.1.0.7.2.9.0-203
	org.apache.atlas	sample-app	2.1.0.7.2.9.0-203
	org.apache.atlas	sqoop-bridge	2.1.0.7.2.9.0-203
	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.2.9.0-203
Apache Avro	org.apache.avro	avro	1.8.2.7.2.9.0-203
	org.apache.avro	avro-compiler	1.8.2.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.avro	avro-ipc	1.8.2.7.2.9.0-203
	org.apache.avro	avro-mapred	1.8.2.7.2.9.0-203
	org.apache.avro	avro-maven-plugin	1.8.2.7.2.9.0-203
	org.apache.avro	avro-protobuf	1.8.2.7.2.9.0-203
	org.apache.avro	avro-service-archetype	1.8.2.7.2.9.0-203
	org.apache.avro	avro-thrift	1.8.2.7.2.9.0-203
	org.apache.avro	avro-tools	1.8.2.7.2.9.0-203
	org.apache.avro	trevni-avro	1.8.2.7.2.9.0-203
	org.apache.avro	trevni-core	1.8.2.7.2.9.0-203
Apache Calcite	org.apache.calcite	calcite-babel	1.21.0.7.2.9.0-203
	org.apache.calcite	calcite-core	1.21.0.7.2.9.0-203
	org.apache.calcite	calcite-druid	1.21.0.7.2.9.0-203
	org.apache.calcite	calcite-kafka	1.21.0.7.2.9.0-203
	org.apache.calcite	calcite-linq4j	1.21.0.7.2.9.0-203
	org.apache.calcite	calcite-server	1.21.0.7.2.9.0-203
	org.apache.calcite	calcite-ubenchmark	1.21.0.7.2.9.0-203
	org.apache.calcite.avatica	avatica	1.16.0.7.2.9.0-203
	org.apache.calcite.avatica	avatica-core	1.16.0.7.2.9.0-203
	org.apache.calcite.avatica	avatica-metrics	1.16.0.7.2.9.0-203
	org.apache.calcite.avatica	avatica-metrics-dropwizardmetrics	1.16.0.7.2.9.0-203
	org.apache.calcite.avatica	avatica-noop-driver	1.16.0.7.2.9.0-203
	org.apache.calcite.avatica	avatica-server	1.16.0.7.2.9.0-203
	org.apache.calcite.avatica	avatica-standalone-server	1.16.0.7.2.9.0-203
	org.apache.calcite.avatica	avatica-tck	1.16.0.7.2.9.0-203
Apache Crunch	org.apache.crunch	crunch-archetype	0.11.0.7.2.9.0-203
	org.apache.crunch	crunch-contrib	0.11.0.7.2.9.0-203
	org.apache.crunch	crunch-core	0.11.0.7.2.9.0-203
	org.apache.crunch	crunch-examples	0.11.0.7.2.9.0-203
	org.apache.crunch	crunch-hbase	0.11.0.7.2.9.0-203
	org.apache.crunch	crunch-hive	0.11.0.7.2.9.0-203
	org.apache.crunch	crunch-scrunch	0.11.0.7.2.9.0-203
	org.apache.crunch	crunch-spark	0.11.0.7.2.9.0-203
	org.apache.crunch	crunch-test	0.11.0.7.2.9.0-203
Apache Druid	org.apache.druid	druid-aws-common	0.17.1.7.2.9.0-203
	org.apache.druid	druid-benchmarks	0.17.1.7.2.9.0-203
	org.apache.druid	druid-console	0.17.1.7.2.9.0-203
	org.apache.druid	druid-core	0.17.1.7.2.9.0-203
	org.apache.druid	druid-gcp-common	0.17.1.7.2.9.0-203
	org.apache.druid	druid-hll	0.17.1.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.druid	druid-indexing-hadoop	0.17.1.7.2.9.0-203
	org.apache.druid	druid-indexing-service	0.17.1.7.2.9.0-203
	org.apache.druid	druid-integration-tests	0.17.1.7.2.9.0-203
	org.apache.druid	druid-processing	0.17.1.7.2.9.0-203
	org.apache.druid	druid-server	0.17.1.7.2.9.0-203
	org.apache.druid	druid-services	0.17.1.7.2.9.0-203
	org.apache.druid	druid-sql	0.17.1.7.2.9.0-203
	org.apache.druid	extendedset	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-hadoop-extensions	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-basic-security	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-bloom-filter	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-datasketches	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-ec2-extensions	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-google-extensions	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-hdfs-storage	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-histogram	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-kafka-extraction-namespace	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-kafka-indexing-service	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-kerberos	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-kinesis-indexing-service	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-lookups-cached-global	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-lookups-cached-single	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-morc-extensions	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-parquet-extensions	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-protobuf-extensions	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-s3-extensions	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-stats	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-sql-metadata-storage	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-sql-metadata-storage	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-ssl-client-sslcontext	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-traffic-emitter	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-traffic-emitter	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-traffic-extensions	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-traffic-storage	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-traffic-extensions	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-traffic-count	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-traffic-extensions	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-traffic-emitter	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-traffic-sketch	0.17.1.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.druid.extensions	druid-scheduler-average-query	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-scheduler-emitter	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-scheduler-cache	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-scheduler-scheduler	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-scheduler-extensions	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-scheduler-min-max	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-scheduler-columns	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-scheduler-contrib	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-scheduler-contrib	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-scheduler-view-maintenance	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-scheduler-view-selection	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-scheduler-data-storage	0.17.1.7.2.9.0-203
	org.apache.druid.extensions	druid-scheduler-contrib	0.17.1.7.2.9.0-203
GCS Connector	com.google.cloud.bigtable	gcs-connector	2.1.2.7.2.9.0-203
	com.google.cloud.bigtable	gcs-connector	2.1.2.7.2.9.0-203
	com.google.cloud.bigtable	gcs-connector	2.1.2.7.2.9.0-203
	com.google.cloud.bigtable	gcs-connector	2.1.2.7.2.9.0-203
	com.google.cloud.bigtable	gcs-connector	2.1.2.7.2.9.0-203
Apache Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-annotations	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-archives	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-assemblies	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-auth	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-aws	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-azure	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-build-tools	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-client	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-client-api	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-client-miniclust	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-common	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-datajoin	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-distcp	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-extras	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-fs2img	3.1.1.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-gridmix	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-client	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-common	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-config	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-container-service	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-docs	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-hadoop-dependency-client	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-hadoop-dependency-server	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-hadoop-dependency-test	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-interface-admin	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-interface-client	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-interface-server	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-server-framework	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-server-scm	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-test-utils	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdds-tools	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-hdfs	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-kafka	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-kms	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-minicluster	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-minikdc	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-nfs	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-openstack	3.1.1.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-ozone-client	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-common	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-csi	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-datanode	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-dist	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-fsfilesystem	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-fsfilesystem-common	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-fsfilesystem-hadoop2	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-fsfilesystem-hadoop3	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-fsfilesystem-shaded	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-insight	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-integration-test	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-interface-client	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-interface-storage	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-network-tests	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-ozone-manager	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-recon	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-reconcodegen	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-s3gateway	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-ozone-tools	1.0.0.7.2.9.0-203
	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-rumen	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-sls	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-streaming	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.2.9.0-203
	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.2.9.0-203
	org.apache.hadoop	mini-chaos-tests	1.0.0.7.2.9.0-203
Apache HBase	org.apache.hbase	hbase-annotations	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-checkstyle	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-client	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-client-project	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-common	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-endpoint	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-examples	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-external-blockcache	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-hadoop-compat	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-hadoop2-compat	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-hbtop	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-http	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-it	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-mapreduce	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-metrics	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-metrics-api	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-procedure	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-protocol	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-protocol-shaded	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-replication	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-resource-bundle	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-rest	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-rsgroup	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-server	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-shaded-client	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-shaded-client-project	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-shaded-mapreduce	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-shaded-testing-util	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-shaded-testing-util-tester	2.2.6.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.hbase	hbase-shell	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-testing-util	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-thrift	2.2.6.7.2.9.0-203
	org.apache.hbase	hbase-zookeeper	2.2.6.7.2.9.0-203
	org.apache.hbase.contrib	hbase-kudu-model	1.0.0.7.2.9.0-203
	org.apache.hbase.contrib	hbase-kudu-proxy	1.0.0.7.2.9.0-203
	org.apache.hbase.contrib	hbase-spark	1.0.0.7.2.9.0-203
	org.apache.hbase.contrib	hbase-spark-it	1.0.0.7.2.9.0-203
	org.apache.hbase.contrib	hbase-spark-protocol	1.0.0.7.2.9.0-203
	org.apache.hbase.contrib	hbase-spark-protocol-shaded	1.0.0.7.2.9.0-203
	org.apache.hbase.file	hbasefs-impl	1.0.0.7.2.9.0-203
	org.apache.hbase.file	hbasefs	1.0.0.7.2.9.0-203
	org.apache.hbase.thrift	hbase-shaded-gson	3.3.0.7.2.9.0-203
	org.apache.hbase.thrift	hbase-shaded-miscellaneous	3.3.0.7.2.9.0-203
	org.apache.hbase.thrift	hbase-shaded-netty	3.3.0.7.2.9.0-203
	org.apache.hbase.thrift	hbase-shaded-protobuf	3.3.0.7.2.9.0-203
Apache Hive	org.apache.hive	hive-accumulo-handler	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-beeline	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-blobstore	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-classification	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-cli	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-common	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-contrib	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-druid-handler	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-exec	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-hbase-handler	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-hplsql	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-it-custom-serde	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-it-druid	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-it-minikdc	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-it-qfile	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-it-qfile-accumulo	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-it-test-serde	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-it-unit	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-it-util	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-jdbc	3.1.3000.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.hive	hive-jdbc-handler	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-jmh	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-kryo-registrator	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-kudu-handler	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-llap-client	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-llap-common	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-llap-ext-client	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-llap-server	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-llap-tez	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-metastore	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-pre-upgrade	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-serde	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-service	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-service-rpc	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-shims	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-spark-client	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-standalone-metastore	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-storage-api	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-streaming	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-testutils	3.1.3000.7.2.9.0-203
	org.apache.hive	hive-vector-code-gen	3.1.3000.7.2.9.0-203
	org.apache.hive	kafka-handler	3.1.3000.7.2.9.0-203
	org.apache.hive.hcatalog	hive-hcatalog-core	3.1.3000.7.2.9.0-203
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	3.1.3000.7.2.9.0-203
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	3.1.3000.7.2.9.0-203
	org.apache.hive.hcatalog	hive-hcatalog-streaming	3.1.3000.7.2.9.0-203
	org.apache.hive.hcatalog	hive-webhcat	3.1.3000.7.2.9.0-203
	org.apache.hive.hcatalog	hive-webhcat-java-client	3.1.3000.7.2.9.0-203
	org.apache.hive.hive-udf-classloader-udf1 it-custom-udfs		3.1.3000.7.2.9.0-203
	org.apache.hive.hive-udf-classloader-udf2 it-custom-udfs		3.1.3000.7.2.9.0-203
	org.apache.hive.hive-udf-classloader-util it-custom-udfs		3.1.3000.7.2.9.0-203
	org.apache.hive.hive-udf-vectorized-badexample it-custom-udfs		3.1.3000.7.2.9.0-203
	org.apache.hive.shims	hive-shims-0.20	3.1.3000.7.2.9.0-203
	org.apache.hive.shims	hive-shims-0.23	3.1.3000.7.2.9.0-203
	org.apache.hive.shims	hive-shims-common	3.1.3000.7.2.9.0-203
	org.apache.hive.shims	hive-shims-scheduler	3.1.3000.7.2.9.0-203

Project	groupId	artifactId	version
Apache Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.2.9.0-203
Apache Kafka	org.apache.kafka	connect	2.5.0.7.2.9.0-203
	org.apache.kafka	connect-api	2.5.0.7.2.9.0-203
	org.apache.kafka	connect-basic-auth-extension	2.5.0.7.2.9.0-203
	org.apache.kafka	connect-file	2.5.0.7.2.9.0-203
	org.apache.kafka	connect-json	2.5.0.7.2.9.0-203
	org.apache.kafka	connect-mirror	2.5.0.7.2.9.0-203
	org.apache.kafka	connect-mirror-client	2.5.0.7.2.9.0-203
	org.apache.kafka	connect-runtime	2.5.0.7.2.9.0-203
	org.apache.kafka	connect-transforms	2.5.0.7.2.9.0-203
	org.apache.kafka	generator	2.5.0.7.2.9.0-203
	org.apache.kafka	jmh-benchmarks	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-clients	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-cloudera-plugins	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-examples	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-log4j-appender	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-examples	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-scala_2.12	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-scala_2.13	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-test-utils	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-upgrade-system-tests-10	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-upgrade-system-tests-11	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-upgrade-system-tests-20	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-upgrade-system-tests-21	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-upgrade-system-tests-22	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-upgrade-system-tests-23	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-streams-upgrade-system-tests-24	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka-tools	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka_2.12	2.5.0.7.2.9.0-203
	org.apache.kafka	kafka_2.13	2.5.0.7.2.9.0-203
Apache Knox	org.apache.knox	gateway-adapter	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-admin-ui	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-applications	1.3.0.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.knox	gateway-cloud-bindings	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-demo-ldap	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-discovery-ambari	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-discovery-cm	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-docker	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-i18n	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-i18n-logging-slf4j	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-performance-test	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-ha	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-jersey	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-rewrite	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-release	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-server	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-server-launcher	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.knox	gateway-service-admin	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-as	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-definitions	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-hbase	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-health	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-hive	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-idbroker	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-impala	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-jkg	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-knoxsso	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-knoxssout	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-livy	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-metadata	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-nifi	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-rm	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-session	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-storm	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-test	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-tgs	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-vault	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-service-webhdfs	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-shell	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-shell-launcher	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-shell-release	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-shell-samples	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-spi	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-test	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-test-idbroker	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-test-release-utils	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-test-utils	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-topology-simple	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-util-common	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-util-configinjector	1.3.0.7.2.9.0-203
	org.apache.knox	gateway-util-launcher	1.3.0.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.knox	gateway-util-urltemplate	1.3.0.7.2.9.0-203
	org.apache.knox	hadoop-examples	1.3.0.7.2.9.0-203
	org.apache.knox	knox-cli-launcher	1.3.0.7.2.9.0-203
	org.apache.knox	knox-homepage-ui	1.3.0.7.2.9.0-203
	org.apache.knox	webhdfs-kerb-test	1.3.0.7.2.9.0-203
	org.apache.knox	webhdfs-test	1.3.0.7.2.9.0-203
Apache Kudu	org.apache.kudu	kudu-backup-tools	1.14.0.7.2.9.0-203
	org.apache.kudu	kudu-backup2_2.11	1.14.0.7.2.9.0-203
	org.apache.kudu	kudu-backup3_2.12	1.14.0.7.2.9.0-203
	org.apache.kudu	kudu-client	1.14.0.7.2.9.0-203
	org.apache.kudu	kudu-hive	1.14.0.7.2.9.0-203
	org.apache.kudu	kudu-spark2-tools_2.11	1.14.0.7.2.9.0-203
	org.apache.kudu	kudu-spark2_2.11	1.14.0.7.2.9.0-203
	org.apache.kudu	kudu-spark3-tools_2.12	1.14.0.7.2.9.0-203
	org.apache.kudu	kudu-spark3_2.12	1.14.0.7.2.9.0-203
	org.apache.kudu	kudu-test-utils	1.14.0.7.2.9.0-203
Apache Livy	org.apache.livy	livy-api	0.6.0.7.2.9.0-203
	org.apache.livy	livy-client-common	0.6.0.7.2.9.0-203
	org.apache.livy	livy-client-http	0.6.0.7.2.9.0-203
	org.apache.livy	livy-core_2.11	0.6.0.7.2.9.0-203
	org.apache.livy	livy-examples	0.6.0.7.2.9.0-203
	org.apache.livy	livy-integration-test	0.6.0.7.2.9.0-203
	org.apache.livy	livy-repl_2.11	0.6.0.7.2.9.0-203
	org.apache.livy	livy-rsc	0.6.0.7.2.9.0-203
	org.apache.livy	livy-scala-api_2.11	0.6.0.7.2.9.0-203
	org.apache.livy	livy-server	0.6.0.7.2.9.0-203
	org.apache.livy	livy-test-lib	0.6.0.7.2.9.0-203
	org.apache.livy	livy-thriftserver	0.6.0.7.2.9.0-203
	org.apache.livy	livy-thriftserver-session	0.6.0.7.2.9.0-203
Apache Lucene	org.apache.lucene	lucene-analyzers-common	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-analyzers-icu	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-analyzers-kuromoji	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-analyzers-morfologik	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-analyzers-nori	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-analyzers-openssl	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-analyzers-phonetic	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-analyzers-smartcn	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-analyzers-stempel	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-backward-codecs	8.4.1.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.lucene	lucene-benchmark	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-classification	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-codecs	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-core	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-demo	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-expressions	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-facet	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-grouping	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-highlighter	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-join	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-memory	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-misc	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-monitor	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-queries	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-queryparser	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-replicator	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-sandbox	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-spatial	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-spatial-extras	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-spatial3d	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-suggest	8.4.1.7.2.9.0-203
	org.apache.lucene	lucene-test-framework	8.4.1.7.2.9.0-203
Apache Oozie	org.apache.oozie	oozie-client	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-core	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-distro	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-examples	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-server	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-sharelib-git	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.2.9.0-203
	org.apache.oozie	oozie-tools	5.1.0.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.2.9.0-203
	org.apache.oozie.test	oozie-mini	5.1.0.7.2.9.0-203
Apache ORC	org.apache.orc	orc-core	1.5.1.7.2.9.0-203
	org.apache.orc	orc-examples	1.5.1.7.2.9.0-203
	org.apache.orc	orc-mapreduce	1.5.1.7.2.9.0-203
	org.apache.orc	orc-shims	1.5.1.7.2.9.0-203
	org.apache.orc	orc-tools	1.5.1.7.2.9.0-203
Apache Parquet	org.apache.parquet	parquet-avro	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-cascading	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-cascading3	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-column	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-common	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-encoding	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-format-structures	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-generator	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-hadoop	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-jackson	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-pig	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-pig-bundle	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-protobuf	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-scala_2.10	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-thrift	1.10.99.7.2.9.0-203
	org.apache.parquet	parquet-tools	1.10.99.7.2.9.0-203
Apache Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.2	5.1.1.7.2.9.0-203
	org.apache.phoenix	phoenix-client-hbase-2.2	5.1.1.7.2.9.0-203
	org.apache.phoenix	phoenix-connectors-phoenix5-compat	6.0.0.7.2.9.0-203
	org.apache.phoenix	phoenix-core	5.1.1.7.2.9.0-203
	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.1.7.2.9.0-203
	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.1.7.2.9.0-203
	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.1.7.2.9.0-203
	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.1.7.2.9.0-203
	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.1.7.2.9.0-203
	org.apache.phoenix	phoenix-pherf	5.1.1.7.2.9.0-203
	org.apache.phoenix	phoenix-queryserver	6.0.0.7.2.9.0-203
	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.2.9.0-203
	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.2.9.0-203
	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.2.9.0-203
	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.phoenix	phoenix-server-hbase-2.2	5.1.1.7.2.9.0-203
	org.apache.phoenix	phoenix-tools	5.1.1.7.2.9.0-203
	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.2.9.0-203
	org.apache.phoenix	phoenix5-hive	6.0.0.7.2.9.0-203
	org.apache.phoenix	phoenix5-spark	6.0.0.7.2.9.0-203
	org.apache.phoenix	phoenix-shaded-commons-cli	1.1.0.7.2.9.0-203
	org.apache.phoenix	phoenix-shaded-guava	1.1.0.7.2.9.0-203
Apache Ranger	org.apache.ranger	conditions-enrichers	2.1.0.7.2.9.0-203
	org.apache.ranger	credentialbuilder	2.1.0.7.2.9.0-203
	org.apache.ranger	embeddedwebserver	2.1.0.7.2.9.0-203
	org.apache.ranger	jisql	2.1.0.7.2.9.0-203
	org.apache.ranger	ldapconfigcheck	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-adls-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-atlas-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-atlas-plugin-shim	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-distro	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-examples-distro	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-hbase-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-hbase-plugin-shim	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-hdfs-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-hdfs-plugin-shim	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-hive-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-hive-plugin-shim	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-intg	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-kafka-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-kafka-plugin-shim	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-kms	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-kms-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-kms-plugin-shim	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-knox-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-knox-plugin-shim	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-kudu-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-kylin-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-kylin-plugin-shim	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-nifi-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-nifi-registry-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-ozone-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-ozone-plugin-shim	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-plugin-classloader	2.1.0.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.ranger	ranger-plugins-audit	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-plugins-common	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-plugins-cred	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-plugins-installer	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-raz-adls	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-raz-hook-abfs	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-raz-hook-s3	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-raz-intg	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-raz-processor	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-raz-s3	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-raz-s3-lib	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-rms-common	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-rms-hive	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-rms-plugins-common	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-rms-webapp	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-s3-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-sampleapp-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-schema-registry-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-solr-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-solr-plugin-shim	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-sqoop-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-sqoop-plugin-shim	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-storm-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-storm-plugin-shim	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-tagsync	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-tools	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-util	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-yarn-plugin	2.1.0.7.2.9.0-203
	org.apache.ranger	ranger-yarn-plugin-shim	2.1.0.7.2.9.0-203
	org.apache.ranger	sample-client	2.1.0.7.2.9.0-203
	org.apache.ranger	sampleapp	2.1.0.7.2.9.0-203
	org.apache.ranger	shaded-raz-hook-abfs	2.1.0.7.2.9.0-203
	org.apache.ranger	shaded-raz-hook-s3	2.1.0.7.2.9.0-203
	org.apache.ranger	ugsync-util	2.1.0.7.2.9.0-203
	org.apache.ranger	unixauthclient	2.1.0.7.2.9.0-203
	org.apache.ranger	unixauthservice	2.1.0.7.2.9.0-203
	org.apache.ranger	unixusersync	2.1.0.7.2.9.0-203
Apache Solr	org.apache.solr	solr-analysis-extras	8.4.1.7.2.9.0-203
	org.apache.solr	solr-analytics	8.4.1.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.solr	solr-cell	8.4.1.7.2.9.0-203
	org.apache.solr	solr-clustering	8.4.1.7.2.9.0-203
	org.apache.solr	solr-core	8.4.1.7.2.9.0-203
	org.apache.solr	solr-dataimporthandler	8.4.1.7.2.9.0-203
	org.apache.solr	solr-dataimporthandler-extras	8.4.1.7.2.9.0-203
	org.apache.solr	solr-jaegertracer-configurator	8.4.1.7.2.9.0-203
	org.apache.solr	solr-langid	8.4.1.7.2.9.0-203
	org.apache.solr	solr-ltr	8.4.1.7.2.9.0-203
	org.apache.solr	solr-prometheus-exporter	8.4.1.7.2.9.0-203
	org.apache.solr	solr-security-util	8.4.1.7.2.9.0-203
	org.apache.solr	solr-solrj	8.4.1.7.2.9.0-203
	org.apache.solr	solr-test-framework	8.4.1.7.2.9.0-203
	org.apache.solr	solr-velocity	8.4.1.7.2.9.0-203
Apache Spark	org.apache.spark	spark-avro_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-catalyst_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-core_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-graphx_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-hadoop-cloud_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-hive_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-kubernetes_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-kvstore_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-launcher_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-mllib-local_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-mllib_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-network-common_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-network-shuffle_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-network-yarn_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-repl_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-shaded-raz	2.4.7.7.2.9.0-203
	org.apache.spark	spark-sketch_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-sql_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-streaming_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-tags_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-unsafe_2.11	2.4.7.7.2.9.0-203
	org.apache.spark	spark-yarn_2.11	2.4.7.7.2.9.0-203

Project	groupId	artifactId	version
Apache Sqoop	org.apache.sqoop	sqoop	1.4.7.7.2.9.0-203
	org.apache.sqoop	sqoop-test	1.4.7.7.2.9.0-203
Apache Tez	org.apache.tez	hadoop-shim	0.9.1.7.2.9.0-203
	org.apache.tez	hadoop-shim-2.8	0.9.1.7.2.9.0-203
	org.apache.tez	tez-api	0.9.1.7.2.9.0-203
	org.apache.tez	tez-aux-services	0.9.1.7.2.9.0-203
	org.apache.tez	tez-common	0.9.1.7.2.9.0-203
	org.apache.tez	tez-dag	0.9.1.7.2.9.0-203
	org.apache.tez	tez-examples	0.9.1.7.2.9.0-203
	org.apache.tez	tez-ext-service-tests	0.9.1.7.2.9.0-203
	org.apache.tez	tez-history-parser	0.9.1.7.2.9.0-203
	org.apache.tez	tez-javadoc-tools	0.9.1.7.2.9.0-203
	org.apache.tez	tez-job-analyzer	0.9.1.7.2.9.0-203
	org.apache.tez	tez-mapreduce	0.9.1.7.2.9.0-203
	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.2.9.0-203
	org.apache.tez	tez-runtime-internals	0.9.1.7.2.9.0-203
	org.apache.tez	tez-runtime-library	0.9.1.7.2.9.0-203
	org.apache.tez	tez-tests	0.9.1.7.2.9.0-203
	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.2.9.0-203
	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.2.9.0-203
	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.2.9.0-203
	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.2.9.0-203
Apache Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.2.9.0-203
	org.apache.zeppelin	zeppelin-display	0.8.2.7.2.9.0-203
	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.2.9.0-203
	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.2.9.0-203
	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.2.9.0-203
	org.apache.zeppelin	zeppelin-livy	0.8.2.7.2.9.0-203
	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.2.9.0-203
	org.apache.zeppelin	zeppelin-server	0.8.2.7.2.9.0-203
	org.apache.zeppelin	zeppelin-shaded-raz	0.8.2.7.2.9.0-203
	org.apache.zeppelin	zeppelin-shell	0.8.2.7.2.9.0-203
	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.2.9.0-203
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.5.5.7.2.9.0-203
	org.apache.zookeeper	zookeeper-client-c	3.5.5.7.2.9.0-203
	org.apache.zookeeper	zookeeper-contrib-loggraph	3.5.5.7.2.9.0-203
	org.apache.zookeeper	zookeeper-contrib-rest	3.5.5.7.2.9.0-203
	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.5.5.7.2.9.0-203
	org.apache.zookeeper	zookeeper-docs	3.5.5.7.2.9.0-203

Project	groupId	artifactId	version
	org.apache.zookeeper	zookeeper-jute	3.5.5.7.2.9.0-203
	org.apache.zookeeper	zookeeper-recipes-election	3.5.5.7.2.9.0-203
	org.apache.zookeeper	zookeeper-recipes-lock	3.5.5.7.2.9.0-203
	org.apache.zookeeper	zookeeper-recipes-queue	3.5.5.7.2.9.0-203

What's New In Cloudera Runtime 7.2.9

You must be aware of the additional functionalities and improvements to features of components in Cloudera Runtime 7.2.9. Learn how the new features and improvements benefit you.

What's New in Apache Hive

There are no new features for Apache Hive in this release of Cloudera Runtime 7.2.9.

What's New in Hue

Learn about the new features of Hue in Cloudera Runtime 7.2.9.

Added Phoenix as a new supported dialect

Apache Phoenix provides a SQL interface to create, access, and query HBase tables using the Hue query editor. You can query HBase tables using the Phoenix SQL Data Definition Language (DDL) and Data Manipulation Language (DML) commands and grammar. Phoenix dialect is available out of the box when the cluster is created using the “Operational Database with SQL” cluster template.

What's New in Apache Impala

Learn about the new features of Impala in Cloudera Runtime 7.2.9.

Although the technical preview feature listed here exists with CDP 7.2.9, it is not yet officially supported by Cloudera. This technical preview feature is still under development and is not recommended for a production environment.

Support GCS Storage (Technical Preview only)

This release added GCS as a storage system for Impala along with the other currently supported storage systems. You can use Impala to query data residing on the GCS. This capability allows convenient access to a storage system that is remotely managed, accessible from anywhere, and integrated with various cloud-based services.

What's New in Apache Kafka

Learn about the new features of Kafka in Cloudera Runtime 7.2.9.

Kafka metrics endpoints can be secured

TLS/SSL encryption as well as Basic Authentication can now be enabled for the HTTP endpoints of Kafka metrics.

TLS/SSL encryption can be enabled with the SSL Encryption For HTTP Metric Reporter property. However, this property does not take effect on its own, TLS/SSL must be enabled for the Kafka broker for this property to work.

Basic Authentication can be enabled with the Enable HTTP Metric Report Basic Authentication property. In addition to enabling Basic Authentication, you also need to set a username and password used for authentication. These can be configured with the HTTP Metric Report User and HTTP Metric Report Password properties respectively.

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

What's New in Apache Kudu

Learn about the new features of Kudu in Cloudera Runtime 7.2.9.

Multiple tablet ids in 'local_replica delete'

The 'local_replica delete' tool allows multiple tablet identifiers to be specified and processed at once. This helps to reduce the overall latency caused by that opening tablet server's metadata takes significant time.

Adding --ignore_nonexistent for 'local_replica delete'

--ignore_nonexistent flag was added for the 'local_replica delete' tool. This makes the real-world scripting scenarios easier if trying to clean up tablet servers of particular tablet replicas.

KuduContext track operations per table

Adds the ability to track operation counts per table. Introduces the MapAccumulator to track these metrics in a single accumulator per operation type.

Support columnar row format in Java client

The setRowDataFormat() method is added to KuduScanner and AsyncKuduScanner. The Java client now supports the columnar RPC format. The format can be set through the setRowDataFormat() method on the KuduScanner.

Check range predicate first while evaluating Bloom filter predicate

Range predicates can be specified along with Bloom filter predicates for the same column. It is more effective to check against range predicates and exit early if the column value is out of bounds compared to computing hash and then looking up the value in Bloom filter.

Arenas for RPC request and response

RPC server side allocates a protobuf Arena for each request. The request RPC and response are allocated from the Arena, ensuring that any sub-messages, strings, repeated fields, and so on, use that Arena for allocation as well. Everything is deleted en-masse when the InboundCall object (which owns the Arena) is destroyed.

New metadata to avoid master when using scan tokens

A new metadata is added to the scan token to allow it to contain all of the metadata required to construct a KuduTable and open a scanner in the clients. This means the GetTableSchema and GetTableLocations RPC calls to the master are no longer required when using the scan token.

New TableMetadataPB, TabletMetadataPB, and authorization token fields were added as optional fields on the token. Additionally a `projected_column_idx` field was added that can be used in place of the `projected_columns`. This significantly reduces the size of the scan token by not duplicating the ColumnSchemaPB that is already in the TableMetadataPB.

Adding the table metadata to the scan token is enabled by default. However, it can be disabled in rare cases where more resiliency to column renaming is desired. It can be disabled in the kudu-spark integration using the kudu.useDriverMetadata property.

RaftConsensus::DumpStatusHtml() does not block Raft consensus activity

kudu::consensus::RaftConsensus::CheckLeadershipAndBindTerm() needs to take the lock to check the term and the Raft role. When many RPCs come in for the same tablet, the contention can hog service threads and cause queue overflows on busy systems. With this improvement, RaftConsensus::DumpStatusHtml() no longer blocks Raft consensus activity and is not blocked by it either.

What's New in Apache Phoenix

Learn about the new features of Phoenix in Cloudera Runtime 7.2.9.

Apache Phoenix 5.1.0

Apache Phoenix version in Cloudera Runtime 7.1.9 is Apache Phoenix 5.1.0. You can read about the new capabilities in this page. For more information about all the fixes and improvements in Apache Phoenix 5.1.0, see the *Release 5.1.0* section in the [Phoenix changelog](#).

Transaction support

X-row, X-table transaction support enables support for complex distributed transactions. Transaction support enables you to run atomic database operations, meaning your database operations must either be completed or terminated.

This two-phase commit capability is implemented using ANSI SQL semantics which is the same as used in other databases like MySQL or PostgreSQL, but this implementation benefits from the near-linear scalability of Apache HBase and Apache Phoenix, and do not have upper-bound on the number of transactions, size of rows/table, or size per table like other databases.

For more information about transaction support, see [Transaction Support](#).

For more information about benchmark results, see the [Bringing Transaction Support to Cloudera Operational Database blog post](#).

Strongly consistent indexing support

Strongly consistent indexing is a three-phase indexing algorithm that ensures that the index table is always in sync with the data table. Newly created Apache Phoenix tables use strongly consistent indexing by default. You can migrate your old Apache Phoenix tables to use strongly consistent indexing.

For more information about strongly consistent indexing, see [Strongly consistent indexing](#).

Support for Apache Phoenix in Hue

Apache Phoenix dialect is available out of the box in Hue when the cluster is created using the “Operational Database with SQL” cluster template.

For more information, see [Using SQL to query HBase from Hue](#).

What's New in Schema Registry

Learn about the new features of Schema Registry in Cloudera Runtime 7.2.9.

Import tool for Confluent Schema Registry

You can now import schemas from Confluent Schema Registry's Kafka topic dump. The schemas need to be uploaded as a file to Schema Registry via REST.

The Ranger model was modified and a new permission for "export-import" was added. When authorization is enabled, users need to have the Create permission under "export-import" in order to import the schemas.

For more information, see *Importing Confluent Schema Registry schemas into Cloudera Schema Registry*.

Related Information

[Importing Confluent Schema Registry schemas into Cloudera Schema Registry](#)

What's New in Apache Spark

Learn about the new features of Spark in Cloudera Runtime 7.2.9.

Data engineering cluster

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

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

What's New in Sqoop

Learn what's new in the Apache Sqoop client in Cloudera Runtime 7.2.9.

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

Discontinued maintenance of direct mode

The Sqoop direct mode feature is no longer maintained. This feature was primarily designed to import data from an abandoned database, which is no longer updated. Using direct mode has several drawbacks:

- Imports can cause an intermittent and overlapping input split.
- Imports can generate duplicate data.
- Many problems, such as intermittent failures, can occur.
- Additional configuration is required.

Do not use the `--direct` option in Sqoop import or export commands.

What's new in Streams Messaging Manager

Learn about the new features of Streams Messaging Manager in Cloudera Runtime 7.2.9.

Added timestamp to SMM UI logs

SMM UI logs have timestamps.

Kafka supports producer metric whitelisting

The feature allows producerMetrics to be whitelisted based on a whitelist property that should contain a regex with the allowed client.ids.

In order to enable the feature, set the following properties within Kafka Broker Advanced Configuration Snippet (Safety Valve) for `kafka.properties/kafka.properties_role_safety_valve`:

- `producer.metrics.whitelist.enabled=true`
- `producer.metrics.whitelist=regex` containing the whitelist for the desired producer's client.ids

This regex uses the simple java `java.util.regex.Pattern` class to compile the regex, and uses the `match()` method on the client.id to determine whether it fits the regex or not.

This feature deletes all the producer timeSeries entities from Service Monitor (CM metricsStore) that do not match the regex.

Also this feature filters the metrics within Kafka itself, so producerMetrics that do not match the regex, would not be reported by the /api/metrics endpoint in the Kafka brokers. Both CM and Kafka use the same Java class and method to filter the metrics.



Note: If the whitelist values start and end with quotation marks, the quotation marks get removed before compiling the value into a Pattern. The same rule applies if the value starts and ends with whitespaces.

What's New in Apache Hadoop YARN

Learn about the new features of Hadoop YARN in Cloudera Runtime 7.2.9.

Enhanced placement rules

A new placement rules evaluation engine is introduced to resolve some previous limitations. It supports a new JSON-based placement rules format. From Cloudera Runtime 7.2.9, the YARN Queue Manager UI supports only this new placement rules format. Legacy mapping rules are automatically converted into this new format when the YARN Queue Manager UI is used for the first time in Cloudera Runtime 7.2.9 or higher version.

With the new placement rules and evaluation engine, you can define an action that should happen if the target queue of the placement rule does not exist or it cannot be created. This action is called the fallback action and it can have the following values: Skip, PlaceDefault and Reject.

Placement rules can now be created and managed in a more user-friendly way in YARN Queue Manager using Placement rule policies. Placement rule policies provide shortcuts for the most common use cases of placement-rules. All legacy Capacity Scheduler mapping-rules are now available as these policies. They provide the same functionality without having to manually write rules. For example, a previous rule of the form `u:%user:%primary_group.%user` can now be achieved using the `primaryGroupUser` policy. If none of the predefined policies solve your use-case, a custom policy option is available that allows you to define the target queue manually.

For more information about placement rules, see [Managing placement rules](#).

Dynamic child-queues in Weight mode

Support for dynamic child-queues is also introduced for the new Weight mode, continuing similar support in the previously existing relative & absolute modes. As before with other modes, dynamic child-queues for Weight mode can be enabled through the YARN Queue Manager UI.

Dynamic child-queues in the new Weight mode have additional enhancements. First, Weight mode supports two-levels of dynamic queues, while other modes only support one-level. And unlike in other modes where a parent-queue cannot have a mix of static and dynamic children, Weight mode supports parent-queues to have both dynamic and static child-queues simultaneously.

For more information, see [Manage dynamic queues](#).

Weight mode for resource allocation

You can now allocate resources to queues using a new mode called Weight mode. This new mode lets you specify the capacity as weights/fractions of total resources instead of already existing ways of specifying it in percentages or in absolute numbers.

For more information about resource allocation, see [Resource allocation overview](#).

Read-only access to Yarn Queue Manager UI for non-admins (Tech Preview)

You can now allow non-admin users to access YARN Queue Manager in a read-only mode. You can either create a new user account with read-only role or use any existing user account with read-only role in Cloudera Manager to access YARN Queue Manager UI. In the read-only access mode, the user can view all the configurations but cannot make any changes to the configurations.

For more information, see [Provide read-only access to Queue Manager UI](#).

Fair sharing intra-queue preemption support



Note: Fair sharing intra-queue preemption is now a fully supported feature.

You can now enable intra-queue preemption for queues that are configured with fairness-based ordering. If the user has configured ordering-policies for queues (`yarn.scheduler.capacity.<queue-path>.ordering-policy`) to be Fair and if this new feature intra-queue preemption is enabled using YARN Queue Manager UI, then all the applications from the same user get a fair amount of resources. Thus, fair-ordering preemption now ensures that each application under a queue gets its fair-share, whether from a single user or several.

While that is for applications from a single user, for resource allocation across users, you can enforce a limit on each user's resource-usage by setting the user limits (`yarn.scheduler.capacity.<queue-path>.user-limit-factor`).

For more information about ordering policies, see [Set Ordering policies within a specific queue](#).

fs2cs conversion utility tool enhancements

The fs2cs conversion utility tool provides better scheduler transition for customers who are upgrading or migrating from one of the CDH releases.

The conversion from Fair Scheduler to Capacity Scheduler using the fs2cs conversion utility is now better through improved placement rules and a new placement rules evaluation engine (see below), a new Weight mode for capacity allocation (see below) with a new and more flexible auto queue creation feature for Capacity Scheduler.s

For more information, see [Fair Scheduler to Capacity Scheduler transition](#).

Unaffected Components in this release

There are no new features for the following components in Cloudera Runtime 7.2.9: Apache Atlas – Data Analytics Studio – Apache HBase – Apache Hadoop HDFS – Apache Hive – Hue – Apache Knox – Apache Oozie – Apache Ranger – Cloudera Search – Apache Solr – Streams Replication Manager – Apache ZooKeeper

There are no new features for the following components in Cloudera Runtime 7.2.9:

- Apache Atlas
- Data Analytics Studio
- Apache HBase
- Apache Hadoop HDFS
- Apache Hive
- Apache Knox
- Apache Oozie
- Apache Ranger
- Cloudera Search
- Apache Solr
- Streams Replication Manager
- Apache ZooKeeper

Fixed Issues In Cloudera Runtime 7.2.9

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.9.

Fixed Issues in Atlas

Review the list of Apache Atlas issues that are resolved in Cloudera Runtime 7.2.9.

OPSAPS-19748: Update hive_db typedef to have managedlocation.

Added "managedLocation" attribute to hive_db entity in Atlas. This issue is now resolved.

OPSAPS-58847: Atlas TLS protocol excludes changed to TLSv1 and TLSv1.1 instead of earlier TLSv1.2

This issue is now resolved.

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

This issue is now resolved.

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

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

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

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

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

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

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

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

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

CDPD-10873

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

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

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

CDPD-11681:

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

CDPD-13199: Incorrect attribute values in bulk import

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

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

A Spark session can include multiple queries. When Atlas reports the Spark metadata, it creates a single process entity to correspond to the Spark session. The result was that an Atlas lineage picture showed multiple input entities or multiple output entities for a process, but the inputs and outputs were only related by the fact that they were included in operations in the same Spark session. In this

release, the Spark Atlas Connector produces a `spark_application` entity for each Spark job. Each data flow produced by the job creates a `spark_process` entity in Atlas, which tracks the actual input and output data sets for that process. For more information, see [Spark metadata collection](#).

CDPD-12620: Migration progress bar not refreshed

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

This issue is resolved.

CDPD-10151: Short Spark job processes may be lost

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

This issue is resolved.

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

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

This issue is resolved.

CDPD-4662: Lineage graph links not working

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

This issue is resolved.

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

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

This issue is resolved.

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

Apache patch information

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

- ATLAS-4088
- ATLAS-4122
- ATLAS-4127
- ATLAS-3913
- ATLAS-4083
- ATLAS-4080
- ATLAS-4077
- ATLAS-4068
- ATLAS-4024
- ATLAS-4051
- ATLAS-4050
- ATLAS-4081
- ATLAS-4062
- ATLAS-4034
- ATLAS-4055
- ATLAS-4013
- ATLAS-4048

- ATLAS-3911
- ATLAS-3941
- ATLAS-3911
- ATLAS-4022
- ATLAS-4026
- ATLAS-3875
- ATLAS-4016
- ATLAS-4025
- ATLAS-4023
- ATLAS-4015
- ATLAS-4006
- ATLAS-4007
- ATLAS-4010
- ATLAS-4005
- ATLAS-4011
- ATLAS-3743
- ATLAS-4008
- ATLAS-3667
- ATLAS-3864
- ATLAS-4170
- ATLAS-4137
- ATLAS-4111
- ATLAS-4092
- ATLAS-4110
- ATLAS-4107
- ATLAS-4019
- ATLAS-2932
- ATLAS-4057
- ATLAS-3988
- ATLAS-4086
- ATLAS-4102
- ATLAS-4101
- ATLAS-4099
- ATLAS-4091
- ATLAS-4095
- ATLAS-4093
- ATLAS-4097
- ATLAS-4017
- ATLAS-4088
- ATLAS-4083
- ATLAS-3980
- ATLAS-4090
- ATLAS-4063
- ATLAS-4089
- ATLAS-4081
- ATLAS-3913
- ATLAS-4080
- ATLAS-4077
- ATLAS-4068
- ATLAS-4024

- ATLAS-4051
- ATLAS-4050
- ATLAS-4073
- ATLAS-4078
- ATLAS-4058
- ATLAS-3864
- ATLAS-4062
- ATLAS-4046
- ATLAS-4163

Fixed Issues in Cloud Connectors

Review the list of Cloud Connectors issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-18287: Significant thread contention while initialising AzureBlobFileSystem.

Reduced the thread contention while initializing filesystems, especially AzureBlobFileSystem or other object stores. The number of filesystems that can be created in parallel is now limited in the option `fs.creation.parallel.count`, default value 64. A smaller value improves the worker thread startup time in processes, especially when many worker threads are attempting to interact with the same object stores. This issue is now resolved.

Apache patch information

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

- HADOOP-17313

Fixed issues in Cruise Control

Review the list of Cruise Control issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-23610: The `topic_configuration` does not fail with NPE when the broker rack id is not configured.

This issue is resolved.

Fixed issues in Data Analytics Studio

There are no fixed issues for Data Analytics Studio in Cloudera Runtime 7.2.9.

Fixed Issues in Apache Hadoop

Review the list of Hadoop issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-23507: HADOOP-16721: Race condition in raw S3 delete & rename in raw S3 breaks hive.

Delays between S3A directory marker deletion and copying in of (large) files underneath a directory during a rename can report the destination directory of hive parallel renames as non-existent. This only surfaces when S3Guard was disabled. This issue is now resolved.

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

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

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

This issue is now resolved.

Apache Patch Information

- HADOOP-17337
- HADOOP-17483
- HADOOP-17480
- HADOOP-17484
- HADOOP-17433

Fixed Issues in HBase

Review the list of HBase issues that are resolved in Cloudera Runtime 7.2.9.

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

The following HBase metrics are available in Cloudera Manager now:

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

Apache patch information

- HBASE-20598
- HBASE-25221 Backport HBASE-24368
- HBASE-25090
- HBASE-25224
- HBASE-24859
- HBASE-25240
- HBASE-25238
- HBASE-25181
- HBASE-20598
- HBASE-25276
- HBASE-25003 Backport HBASE-24350 and HBASE-24779
- HBASE-25267
- HBASE-25255
- HBASE-25261
- HBASE-23364
- HBASE-25300
- HBASE-25306
- HBASE-25311
- HBASE-25187
- HBASE-25323
- HBASE-24827
- HBASE-24872
- HBASE-25237

- HBASE-25307
- HBASE-25339
- HBASE-25321
- HBASE-25330
- HBASE-25332
- HBASE-25263
- HBASE-23202
- HBASE-25536
- HBASE-24885
- HBASE-25684
- HBASE-25646: Possible Resource Leak in CatalogJanitor
- HBASE-25656 Backport to branch-2.2: [HBASE-25548 Optionally allow snapshots to preserve cluster's max filesize config by setting it into table descriptor]
- HBASE-25636 Expose HBCK report as metrics
- HBASE-25621 Balancer should check region plan source to avoid misplace region groups
- HBASE-25587 [hbck2] Schedule SCP for all unknown servers

Fixed Issues in HDFS

Review the list of HDFS issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-6100

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

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

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

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

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

Apache Patch Information

- HDFS-15632

Fixed Issues in Apache Hive

Review the list of Hive issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-18619: Backport HIVE-24293: Integer overflow in llap collision mask.

This is a fix for a race condition in the cache collision logic. This issue was discovered in DWX-3963 This issue is now resolved.

CDPD-16986: hive.table("table_name").show results in null pointer exception.

Dataframe obtained by hive.table("table_name") is usable now. This issue is now resolved.

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

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

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

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

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

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

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

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

Apache Patch Information

- HIVE-24603
- HIVE-24550
- HIVE-24630
- HIVE-24581
- HIVE-24535
- HIVE-24602
- HIVE-24233
- HIVE-24276
- HIVE-23799
- HIVE-23791
- HIVE-23764
- HIVE-23107
- HIVE-20137

Technical Service Bulletins

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

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

Fixed Issues in Hue

There are no fixed issues for Hue in Cloudera Runtime 7.2.9.

Fixed Issues in Apache Impala

Review the list of Impala issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-23692: Deadloop in table metadata loading when using an invalid RemoteIterator.

Fixed the bug that was causing catalogd to come into an infinite loop when the underlying directories of a table/partition is absent in the filesystem. This only exists on cloud filesystems like S3, ABFS. For more information, see IMPALA-10579. This issue is now resolved.

CDPD-20456: Upgrade jackson to 2.10.5.1 or 2.11.0+ due to CVE-2020-25649.

The fix upgraded jackson to 2.10.5.1 for Impala and it was merged into 7.2.8.x branch and CDH-7.1-maint branch. This issue is now resolved.

CDPD-19304: Upgrade to slf4j 1.7.30.

The fix upgraded slf4j to 1.7.30 for Impala and merged into 7.2.9.x branch and CDH-7.1-maint branch. This issue is now resolved.

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

This issue is now resolved.

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

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

Technical Service Bulletins

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

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

Fixed Issues in Apache Kafka

There are no fixed issues for Kafka in Cloudera Runtime 7.2.9.

Fixed Issues in Apache Knox

Review the list of Knox issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-22677: Fetching cloud token failed with 400.

This fix ensures that IDbroker (Azure) will retry when there is a 400 failure from Azure backend. This issue is now resolved.

CDPD-23016: JSESSIONID cookie missing when Zeppelin UI proxied via Knox.

This fix ensures that set-cookie header attribute order is preserved. This issue is now resolved.

CDPD-5541: Load balancing mode for HS2 and other services.

This feature enables services that use Knox HA provider to loadbalance between backend services with an ability to turn on sticky sessions. This issue is now resolved.

CDPD-19999: Incorrect URLs are produced for failover when you access the NiFi UI in data hub.

This issue is now resolved.

CDPD-20187: Knox HA Dispatch is unable to mark a host as failed without retry.

The services can be configured to failover without a retry. This issue is now resolved.

CDPD-20188: When accessing the Cloudera Manager UI through a Knox proxy, the add Service Wizard fails at the Assign Roles step with message that "A server error has occurred; affects CDH 7.1.5, all Cloudera Manager versions.

This issue is now resolved.

CDPD-20684: Knox loadbalancing for Ranger component is not supported as expected.

Knox loadbalancing for Ranger component now supports the Knox HA Provider parameters `enableStickySession` and `enableLoadBalancing`. This issue is now resolved.

CDPD-19110: Prevent Knox from passing `hadoop.auth` cookie to browser.

This issue is now resolved.

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

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

Apache patch information

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

- KNOX-2511
- KNOX-2406 Use dependency bom for dependency management
- KNOX-2392 Simple file-based TokenStateService implementation

- KNOX-2389 AliasBasedTokenStateService stops processing persisted journal entries if one is malformed
- KNOX-2377 Address potential loss of token state
- KNOX-2384 Token Service should return expiration from token when renewal disabled
- KNOX-2381 racking UI of flink session is broken in YARNUIV2
- KNOX-2378 AliasBasedTokenStateService log message is misleading
- KNOX-2376 Ensure all HBASEJARS IN rules are for /hbase/jars and not /hbase/maven
- KNOX-2368 CM Cluster Configuration Monitor Does Not Support Rolling Restart Events
- KNOX-2351 Catching any errors while monitoring CM configuration changes
- KNOX-2367 Fix rewrite rules for HDFS UI fonts and bootstrap.min.css.map
- KNOX-2348 Fix KnoxCLI when kerberos auth is used
- KNOX-2357 Descriptor handler should not default discovery type to Ambari unless there is discovery configuration
- ODX-107 KNOX-2354 An HBASEJARS service which can proxy HBase jars hosted by tà€
- KNOX-1998 WebHDFS rewrite.xml does not have rewrite rule for Location field in json
- KNOX-2352 Knox Token State Eviction Should Be Based on Expiration and Extended Default Grace Period
- KNOX-2355:Update Atlas rewrite.xml for new UI changes
- KNOX-2304 CM discovery cluster config monitor needs to be aware of tà€
- KNOX-2316 Knox Token State Eviction Must Consider Maximum Token Lifetime
- KNOX-2314 NPE from topology Service equals implementation when no URLs
- KNOX-2301 and KNOX-2302 Trigger discovery for descriptors at gateway start time
- KNOX-2287 KnoxCLI convert topology to provider and descriptor
- KNOX-2298 ClouderaManager cluster config monitor should stop monitoring unreferenced clusters
- KNOX-2266 Tokens Should Include a Unique Identifier
- KNOX-2212 Token permissiveness validation
- KNOX-2230 Token State Service should throw UnknownTokenException instead of IllegalArgumentException
- KNOX-2237 CM service discovery should default the http path of Hive URLs when the associated property is not set
- KNOX-2233 DefaultKeystoreService getCredentialForCluster uses cache without synchronization
- KNOX-2214 Reaping of expired Knox tokens
- KNOX-2228 JWTFilter should handle unknown token exception from token state service
- KNOX-2210 Gateway-level configuration for server-managed Knox token state
- KNOX-2215 Token service should return a 403 response when the renewer is not white-listed
- KNOX-2209 Improve logging for Knox token handling
- KNOX-2153 CM discovery tà€ Monitor Cloudera Manager
- KNOX-2156 CM discovery tà€ KUDUUI discovery
- KNOX-2152 Disable Ambari cluster configuration monitoring by default
- KNOX-2151 HIVE_ON_TEZ HS2 Discovery doesn't work
- KNOX-1970 CM discovery tà€ Add Impala HS2 to auto discovery
- KNOX-1932 CM discovery tà€ WEBHCAT URLs not discovered
- KNOX-1921 CM discovery tà€ Hue Load balancer HTTP/HTTPS scheme
- KNOX-1935 CM discovery tà€ Hue should not have both LB and non LB
- KNOX-1962 CM discovery tà€ Avoid reading krb5 login config multiple tà€
- KNOX-2144 Alias API KnoxShell support should provide response types better than raw JSON
- KNOX-1410 Knox Shell support for remote Alias management
- KNOX-2127 ZooKeeperAliasService mishandles mixed-case alias keys properly
- KNOX-2105 KnoxShell support for token renewal and revocation
- KNOX-2071 Configurable maximum token lifetime for TokenStateService
- KNOX-2066 Composite Authz Provider
- KNOX-2067 KnoxToken service support for renewal and revocation
- KNOX-843 - Add support for load balancing multiple clients across backend service instances

- KNOX-2456 SHS links sometimes broken on FINISHED jobs page (#375) Change-Id: I9d269cd3ed0369d0dc13d0eba8b53bd2da8b1e34
- KNOX-2533 - Qualifying service params for discovery improvements (#401)
- KNOX-2530 - Support qualifying service params for CM discovery control (#398)
- KNOX-2406 - Use dependency bom for dependency management (#363)
- KNOX-2392 - Simple file-based TokenStateService implementation (#350)
- KNOX-2389 - AliasBasedTokenStateService stops processing persisted journal entries if one is malformed (#346)
- KNOX-2377 - Address potential loss of token state (#345)
- KNOX-2384 - Token Service should return expiration from token when renewal disabled (#342)
- KNOX-2381 racking UI of flink session is broken in YARNUIV2 (#340)
- KNOX-2378 - AliasBasedTokenStateService log message is misleading (#339)
- KNOX-2376 Ensure all HBASEJARS IN rules are for /hbase/jars and not /hbase/maven (#338)
- KNOX-2368 - CM Cluster Configuration Monitor Does Not Support Rolling Restart Events
- KNOX-2351 - Catching any errors while monitoring CM configuration changes (#324)
- KNOX-2367 - Fix rewrite rules for HDFS UI fonts and bootstrap.min.css.map (#332)
- KNOX-2348 - Fix KnoxCLI when kerberos auth is used (#331)
- KNOX-2357 - Descriptor handler should not default discovery type to Ambari unless there is discovery configuration (#326)
- KNOX-1998 - WebHDFS rewrite.xml does not have rewrite rule for Location field in json (#138)
- KNOX-2352 - Knox Token State Eviction Should Be Based on Expiration and Extended Default Grace Period (#321)
- KNOX-2355:Update Atlas rewrite.xml for new UI changes
- KNOX-2304 - CM discovery cluster config monitor needs to be aware of ' (€) (#307)
- KNOX-2316 - Knox Token State Eviction Must Consider Maximum Token Lifetime (#306)
- KNOX-2314 - NPE from topology Service equals implementation when no URLs (#303)
- KNOX-2301 and KNOX-2302 (#297)
- KNOX-2287 KnoxCLI convert topology to provider and descriptor
- KNOX-2298 - ClouderaManager cluster config monitor should stop monitoring unreferenced clusters (#291)
- KNOX-2266 - Tokens Should Include a Unique Identifier (#284)
- KNOX-2212 - Token permissiveness validation
- KNOX-2230 - Token State Service should throw UnknownTokenException instead of IllegalArgumentException (#268)
- KNOX-2237 - CM service discovery should default the http path of Hive URLs when the associated property is not set (#266)
- KNOX-2233 - DefaultKeystoreService getCredentialForCluster uses cache without synchronization (#264)
- KNOX-2214 - Reaping of expired Knox tokens
- KNOX-2228 - JWTFilter should handle unknown token exception from token state service (#260)
- KNOX-2210 - Gateway-level configuration for server-managed Knox token state (#259)
- KNOX-2215 - Token service should return a 403 response when the renewer is not white-listed (#251)
- KNOX-2209 - Improve logging for Knox token handling (#250)
- KNOX-2153 - CM discovery - Monitor Cloudera Manager (#239)
- KNOX-2156 - CM discovery - KUDUI discovery (#228)
- KNOX-2152 - Disable Ambari cluster configuration monitoring by default (#225)
- KNOX-2151 - HIVE_ON_TEZ HS2 Discovery doesn't work (#224)
- KNOX-1970 - CM discovery - Add Impala HS2 to auto discovery (#223)
- KNOX-1932 - CM discovery - WEBHCAT URLs not discovered (#222)
- KNOX-1921 - CM discovery - Hue Load balancer HTTP/HTTPS scheme (#221)
- KNOX-1935 - CM discovery - Hue should not have both LB and non LB (#220)
- KNOX-1962 - CM discovery - Avoid reading krb5 login config multiple times (€) (#215)
- KNOX-2144 - Alias API KnoxShell support should provide response types better than raw JSON (#211)

- KNOX-1410 - Knox Shell support for remote Alias management (#210)
- KNOX-2127 - ZooKeeperAliasService mishandles mixed-case alias keys properly (#202)
- KNOX-2105 - KnoxShell support for token renewal and revocation (#180)
- KNOX-2071 - Configurable maximum token lifetime for TokenStateService (#178)
- KNOX-2066 - Composite Authz Provider
- KNOX-2067 - KnoxToken service support for renewal and revocation
- KNOX-2575
- KNOX-2571
- KNOX-2553
- KNOX-2555
- KNOX-2557
- KNOX-2545
- KNOX-2538

Fixed Issues in Apache Kudu

Review the list of Kudu issues that are resolved in Cloudera Runtime 7.2.9.

add cdhVersion field to MDL files

Adds the 'compatibility' and 'cdhVersion' fields to every MDL file, allowing Cloudera Manager to properly detect staleness after a Cloudera Manager upgrade. Previously, Cloudera Manager would unnecessarily list Kudu as a stale service post upgrade, even if the CDH version did not change.

Auto rebalancer does not successfully run moves

When executing moves, the auto-rebalancer used try to resolve the leader's address by passing its UUID instead of its host. With this fix, it uses an appropriate host.

KuduPredicate class in Java client does not handle Date columns

Prior to this fix, if you had a table with DATE column, you could not scan for it using the java client. A check for minimum and maximum boundaries of integer representation of java.sql.Date was added to match MIN_DATE_VALUE and MAX_DATE_VALUE in DateUtil.

- [KUDU-2884](#): Improves the master address matching in the `kudu hms fix` tool
- [KUDU-3142](#): Remove the deprecated kudu-mapreduce modules
- [KUDU-3149](#): Lock contention between registering ops and computing maintenance op stats
- [KUDU-3157](#): Ensure slf4j classes are not shaded
- [KUDU-3182](#): 'last_known_addr' is not specified for single master Raft configuration
- [KUDU-3191](#): Fail replicas when KUDU-2233 is detected
- [KUDU-3195](#): Flush when any DMS in the tablet is older than the time threshold
- [KUDU-3198](#): Fix encodeRow() when encoding delete operations
- [KUDU-3213](#): try at different server on TABLET_NOT_RUNNING
- [KUDU-3254](#): fix bug in meta-cache exposed by KUDU-1802

Apache patch information

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

- KUDU-3090
- KUDU-3091
- KUDU-1563
- KUDU-2966
- KUDU-3210

Fixed Issues in Apache Oozie

Review the list of Oozie issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-20444: The Sqoop build no longer shades the Avro and Parquet libraries as it wasn't needed for a long time.

Oozie now automatically pulls in the necessary avro and parquet libraries into Oozie's Sqoop sharelib. If Avro or Parquet is used with Sqoop with an Oozie Sqoop action then you need not copy these libraries to sharelib manually. This issue is now resolved.

CDPD-21032: Memory leak in EL evaluation.

This issue is now resolved.

CDPD-19684: Oozie Spark action must support automatically copying the hive-site.xml

Oozie will now automatically pick-up the hive-site.xml and add it to the Yarn container of a Spark action. From now on it is not necessary to put a hive-site.xml manually onto Oozie Spark's sharelib. This issue is resolved.

CDPD-21870: Fix Oozie client always using the current system username instead the one specified by the user. For example, through kerberos or explicit basic authentication.

A bug in Oozie CLI caused the Workflow to be launched in the name of the current Unix user even if Kerberos authentication was used with a ticket for a different user. This issue is resolved.

CDPD-23141: No Sqoop logs present in Oozie Sqoop action launcher logs.

Sqoop logs are not present in the aggregated Yarn logs. This issue is now resolved.

CDPD-20002: When stopped, SSH action should stop the spawned processes on target Host if specified.

When an SSH action was killed the child processes launched by the actions were not killed. The default behaviour is still these not getting killed but we introduced 2 ways to do so:

- Use the new 0.3 schema version for your SSH action in your workflow.xml and add the "terminate-subprocesses" XML element with value "true". Example: `<terminate-subprocesses>true</terminate-subprocesses>`
- You can set this globally by adding the following oozie-site.xml safety-valve in Cloudera Manager with value "true" : "oozie.action.ssh.action.terminate.subprocesses"

If both are set then the value set in the workflow.xml takes precedence. This issue is now resolved.

CDPD-20984: Make Oozie backward compatible with the Falcon Oozie EL extension library.

The falcon-oozie-el-extension library written for HDP 2.6 was not compatible with CDP 7.x Oozie. We introduced a change in Oozie to make that library forward compatible with Oozie in CDP 7. NOTE: The rest of the HDP 2.6 Falcon library is still not compatible with CDP but only falcon-oozie-el-extension is. This issue is now resolved.

CDPD-22161: Yarn application should be submitted on code level with a user running the Workflow to avoid asking for a delegation from IdBroker in the name of Oozie.

When the Yarn remote log dir was set to s3 or abfs, log aggregation for Oozie actions was not working by default. The workaround for this was to extend the IdBroker mapping and add the oozie user there, or to add an explicit file-system credential (which pointed to the Yarn remote log dir) to the Oozie Workflow. These workarounds are no longer required as from now on a delegation token for the remote log dir will be obtained from IdBroker in the name of the user who is running the Workflow.. This issue is now resolved.

CDPD-11965: Cookie without HttpOnly and Secure flag set.

The Secure and HttpOnly attributes are now set on all Cookies returned by Oozie as per recommendations. This issue is now resolved.

CDPD-19281: Missing CSP, X-XSS-Protection, HSTS Headers.

Oozie was enhanced with extra HTTP Headers to make it more secure. In scope of these enhancements the following HTTP Headers are now returned by Oozie: X-XSS-Protection with

value "1; mode=block" ; Content-Security-Policy with value "default-src 'self'" ; Strict-Transport-Security with value "max-age=31536000; includeSubDomains".

You can remove these Headers by adding an oozie-site.xml safety-valve with an empty value - should be a simple space - in Cloudera Manager with the "oozie.servlets.response.header." prefix. Example: "oozie.servlets.response.header.Strict-Transport-Security= "

You can also modify the value of these Header the same way through a safety-valve. Example: "oozie.servlets.response.header.Strict-Transport-Security=max-age=604800; includeSubDomains"

Using the same prefix you can also make Oozie return custom HTTP Headers. Example: "oozie.servlets.response.header.MyHeader=MyValue".

These were originally decommissioned when Oozie was rebased from 4.x to 5.x, but to reduce the migration effort for users these are supported again.

This issue is now resolved.

CDPD-19473: CVE-2020-35451 - Fix privilege escalation vulnerability in OozieSharelibCLI.

The security vulnerability was in Oozie's sharelib CLI regarding the temporary directory. This issue is now resolved.

CDPD-20649: Revise yarn.app and mapreduce property overrides in Oozie.

When upgrading from CDH 5 or HDP 2/3 Oozie will still be able to handle the following map-reduce related properties:

- oozie.launcher.mapreduce.map.memory.mb
- oozie.launcher.mapreduce.map.cpu.vcores
- oozie.launcher.mapreduce.map.java.opts

These were originally decommissioned when Oozie was rebased from 4.x to 5.x, but to reduce the migration effort for users these are supported again. This issue is now resolved.

CDPD-21031: Workflow and coordinator action status remains RUNNING after rerun.

This issue is now resolved.

CDPD-18931: No appropriate protocol" error with email action(disable TLS1.0/1.1).

Oozie security enhancements. This issue is now resolved.

CDPD-17598: Log library issues.

The available logging libraries for Hive 2, Spark, and Sqoop actions are adjusted. This issue is now resolved.

CDPD-18703: The Oozie version returns incorrect values.

The "oozie version" command now returns the correct Oozie version and build time. This issue is now resolved.

CDPD-17306: Hive-Common is added as a dependency to Sqoop and Oozie's Sqoop sharelib so that you do not have to do it manually.

This issue is now resolved.

CDPD-17843: Hive-JDBC is added as a dependency to Sqoop and Oozie's Sqoop sharelib so you do not have to do it manually.

This issue is now resolved.

OPSAPS-58298: Oozie must accept a keyStoreType and trustStoreType property in oozie-site.xml.

This issue is now resolved.

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

This issue is now resolved.

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

This issue is now resolved.

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

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

CDPD-14600

Apache ActiveMQ is updated to address CVE-2016-3088

CDPD-13702

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

CDPD-11967

Fix to address CWE-693: Protection Mechanism Failure

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

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

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

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

CDPD-10746: Fix to address CVE-2019-17571**CDPD-9895: Various errors when trying to use an S3 filesystem**

Oozie is now fully compatible with S3.

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

These actions inside the fork action run in parallel mode, and they have some seconds delay in between them. If a parameter is passed to one of these actions, that cannot be resolved, then it changes its status to FAILED, and also the workflow's state to FAILED. The other actions state which are not started yet will stuck in PREP state forever. The correct behaviour would be to stop the remaining actions as well as the workflow. Note: this bug only occurs when it is run in independent mode. If it has a parent workflow, then the parent workflow will stop this workflow after 10 minutes because of the callback process.

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

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

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

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

CDPD-9189: Apache Pig support was completely removed from Oozie**CDPD-7108: In case we have a workflow which has, let's say, 80 actions after each other, then the validator code "never" finishes****CDPD-7107: The following were added to the spark opts section of the spark action: --conf spark****CDPD-7106: query tag is not functional for Hive2 action node in oozie**

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

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

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

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

CDPD-5168: Logging enhancements in CoordElFunctions for better supportability

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

This issue is now fixed.

Apache patch information

- OOZIE-3365
- OOZIE-3596
- OOZIE-3409

Fixed Issues in Phoenix

Review the list of Phoenix issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-21403: Phoenix Omid - Remove netty 3 dependency.

Removed netty 3.x dependency, which is vulnerable to a bunch of CVEs. This issue is resolved.

Apache Patch Information

- PHOENIX-6400
- OMID-200
- OMID-202

Fixed Issues in Parquet

Review the list of Parquet issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-21779: Backport PARQUET-1928.

It is now available to read INT96 values from Parquet files using the parquet-avro binding.

Apache Patch Information

- PARQUET-1493
- PARQUET-1928

Fixed Issues in Apache Ranger

Review the list of Ranger issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-14423: Access audits page not loading.

Fixes Ranger's connection to Solr to pull audit events from Solr Service. Ranger was unable to fetch Audit events from solr after expiry of kerberos ticket. This issue is resolved.

CDPD-17944: Audit-to-cloud storage: minimize write calls.

Added ability to batch the cloud storage write calls. This issue is resolved.

CDPD-20524: Added a capability to specify audit filters from UI side.

This issue is resolved.

CDPD-22820: Handling of invalid usernames for usersync.

Added validation for user/group names to check for invalid characters in usersync before updating to ranger admin. This issue is resolved.

CDPD-23572: Ranger usersync does not synchronise groups.

Allow the setting ranger.usersync.group.searchenabled to false and configure ranger.usersync.ldap.user.groupnameattribute=memberof. That way, usersync can sync the users based on the user search base and user search filter and use the "memberof" attribute of the user to sync all the groups each user belongs to. This issue is resolved.

CDPD-23579: Policy Item does not render in the report page.

This issue is resolved.

CDPD-23726: In place policy or tag updates are by default set to false to resolve performance issue.

This issue is resolved.

CDPD-24387: Ranger Audit framework change to handle UnsupportedOperationException while writing into S3AFileSystem with hflush api.

This issue is resolved.

OPSAPS-14423: Access audits page not loading.

Fixes Ranger's connection to Solr to pull audit events from Solr Service. Ranger was unable to fetch Audit events from solr after expiry of kerberos ticket. This issue is resolved.

OPSAPS-17016: Ranger KMS - Upgrade api-i18n due to CVE-2018-1337.

This issue is resolved.

OPSAPS-13664: RangerRazClient communicates with k5b for every request for populating headers.

Raz performance fix - using connection pool. This issue is resolved.

OPSAPS-13595: Reduce SSL handshake and krb negotiations from RangerRazClient --> Raz Server.

RAZ performance fix - Added Apache Http Client with connection pooling changes. This issue is resolved.

OPSAPS-17467: Upgrade Tomcat from 7.0.x line.

Tomcat is upgraded to 8.5.61. This issue is resolved.

OPSAPS-18273: Upgrade to TLS to version 1.2 and above.

Disabled TLS versions that are less than 1.2 for Ranger. This issue is resolved.

OPSAPS-19638: Thread contention inside AuditFilter.audit() while logging.

Auditing level changes to avoid thread contention due to high volume of auditing. This issue is resolved.

OPSAPS-21704: Ranger Auditor role (API compatibility).

Fixed access for servicedef GET API. This issue is resolved.

OPSAPS-21769: Embedded server max connection configurable.

RAZ performance fix - make connection parameters configurable. This issue is resolved.

OPSAPS-22201: NoClassDefFoundError in Atlas during ranger audit.

Fixed Atlas audit issue by adding right dependency. This issue is resolved.

OPSAPS-22353: Raz for adls is displaying exceptions while executing spark benchmarks.

RAZ performance fix - acceptCount parameter. This issue is resolved.

OPSAPS-23590: Incorrect message when user does not have permission on the storage account.

Fixed the error message returned to have correct "Permission denied" message. This issue is resolved.

OPSAPS-23853: Ranger Raz client blocks waiting for Http connection due to connection leak.

Fixed RAZ client connection leak during failures causing Oozie to not able to get connections. This issue is resolved.

OPSAPS-23917: Unable to access bucket on 7.2.9 RAZ-S3 DL cluster.

Added null check when adding evaluators. This issue is resolved.

CDPD-16888: Ranger and Atlas services should have recommended heap size configured while deploying a cluster.

Default minimum heap size for Ranger services is now set as 1 GB and for Atlas 2GB. This issue is now resolved.

CDPD-16888: Solr client connection used for communication is not closed and this results in resource leak.

This issue is now resolved.

OPSAPS-58711: ODB cannot deploy against Datalake HA

This issue is now resolved.

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

This issue is now resolved. You must use this step as a solution: sudo
ln -s /opt/cloudera/parcels/*<CDH-version>*/jars/libfb303-0.9.3.jar
/opt/cloudera/parcels/*<CDH-version>*/lib/ranger-admin/ews/webapp/WEB-INF/lib/libfb303-0.9.3.jar

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

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

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

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

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

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

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

Apache Patch Information

- RANGER-3208
- RANGER-3189
- RANGER-3153
- RANGER-3168
- RANGER-3202
- RANGER-3194
- RANGER-3147
- RANGER-3226
- RANGER-3213
- RANGER-3209
- RANGER-3205
- RANGER-3203
- RANGER-3210
- RANGER-3199
- RANGER-3189

- RANGER-3207
- RANGER-3191
- RANGER-3194

Fixed Issues in Schema Registry

Review the list of Schema Registry issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-21903: Aggregated schema search does not work correctly

You can now search by schema name and schema description.

Previously, if they tried to search by description only, they would incorrectly get back all schemas in the result set.

CDPD-23028: Delete model-registry module

Deleted "model-registry" from the project as it was not upto-date and not used. Removing it allows Cloudera to resolve several other issues.

CDPD-23183: Introduce Guice and make Atlas an optional feature

In this release, the framework is replaced by Guice. As a consequence, you can no longer add custom modules in the registry.yaml configuration file. Instead, you can programmatically update by changing the java code.

Previously, Schema Registry used a custom-built dependency injection framework.

CDPD-23996: No authorization data in logs

The authorization data for users no longer appears in logs.

OPSAPS-59229: Schema Registry YAML configuration contains plaintext passwords

Passwords are now assigned as environment variables.

Previously, Schema Registry passwords used to be stored as plaintext in the configuration file.

Fixed Issues in Cloudera Search

Review the list of Cloudera Search issues that are resolved in Cloudera Runtime 7.2.9.

SOLR has issues if impersonating principal has a hyphen in its name

Any Kerberos principal which contains characters that do not match regex "[a-zA-Z_][a-zA-Z0-9_]*" cannot be used as an impersonating user in Solr. This may affect HUE and KNOX access to Solr, as the principals for these services are configured as Solr proxy users by default. This issue is resolved.

Fixed Issues in Spark

Review the list of Spark issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-21614: Spark SQL TRUNCATE table not permitted on external purge tables.

In order to retain the legacy Hive1/Hive2 behavior around managed non-acid tables, the migration process instructed to convert those tables to external with external.table.purge=true table property. There were issues that the TRUNCATE TABLE operation cannot be performed through Spark SQL on those tables. Spark now allows you to TRUNCATE an external table if external.table.purge is set to true in table properties. This issue is now resolved.

CDPD-18938: Jobs disappear intermittently from the SHS under high load.

SPARK-33841 has been back-ported to CDPD in order to fix the issue with jobs disappearing intermittently from the SHS under high load. This issue is now resolved.

CDPD-20434: SHS should be resilient to corrupted event log directories.

SPARK-33146 has been back-ported to CDPD in order to make SHS resilient to corrupted event log directories. This issue is now resolved.

CDPD-16010: Removed netty3 dependency.

This replaces an internal patch of Spark Machine Learning events to the community based one. This issue is now resolved.

CDPD-18652: Adapt SAC to new Machine Learning event listener in CDP Spark 2.4

This replaces an internal patch of Spark Machine Learning events to the community based one. This issue is now resolved.

CDPD-16748: Improve LeftSemi SortMergeJoin right side buffering.

This issue is now resolved.

CDPD-17422: Improve null-safe equi-join key extraction.

This issue is now resolved.

CDPD-18458: When pyspark.sql.functions.lit() function is used with dataframe cache, it returns wrong result.

This issue is now resolved.

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

This issue is now resolved.

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

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

This issue is now resolved.

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

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

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

This issue is now resolved.

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

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

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

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

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

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

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

This issue is now resolved.

Apache patch information

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

- SPARK-17875
- SPARK-33841

Fixed Issues in Apache Sqoop

Review the list of Sqoop issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-20444: The Sqoop build no longer shades the Avro and Parquet libraries as it wasn't needed for a long time.

Oozie now automatically pulls in the necessary avro and parquet libraries into Oozie's Sqoop sharelib. If Avro or Parquet is used with Sqoop with an Oozie Sqoop action then you need not copy these libraries to sharelib manually. This issue is now resolved.

CDPD-23157: Sqoop Teradata import fails if source table is empty.

This issue is now resolved.

CDPD-19934: Sqoop should handle keeping custom environment variables when executing beeline in a new process

Until now during Hive import when no `--hs2-url` parameter was specified and Sqoop did the import via the beeline command line utility, in the beeline process launched by Sqoop only the following environment variables were preserved from the parent process: `HADOOP_COMMON_HOME`, `HADOOP_HOME`, `HADOOP_MAPRED_HOME`, `HBASE_HOME`, `HCAT_HOME`, `HIVE_HOME`, `JAVA_HOME`, `PATH`, `ZOOKEEPER_HOME`. Now you can specify custom environment variables to preserve via:

- By specify the "sqoop.beeline.env.preserve" system-property for the Sqoop command. E.g.:
`sqoop import -Dsqoop.beeline.env.preserve=MY_VARIBALE`
- Or by specify the "sqoop.beeline.env.preserve" property as a sqoop-site.xml safety-valve in Cloudera Manager and then it will be applied for every sqoop Hive import.

This issue is now resolved.

CDPD-21148: Remove Accumulo dependency from Sqoop

Accumulo support is removed from Sqoop. This issue is now resolved.

CDPD-18796: In Cloudera Manager, if `sqoop.avro.logical_types.decimal.default.precision` and `sqoop.avro.logical_types.decimal.default.scale` are not set correctly for the Sqoop component, the import job fails with the error "Error: java.lang.ArithmeticException: Rounding necessary".

This issue is now resolved.

CDPD-19647: When the `--query` parameter is used with `--as-orcfile`, Sqoop displays an error.

This issue is now resolved.

CDPD-17306: Hive-Common is added as a dependency to Sqoop and Oozie's Sqoop sharelib so that you do not have to do it manually.

This issue is now resolved.

CDPD-17843: Hive-JDBC is added as a dependency to Sqoop and Oozie's Sqoop sharelib so you do not have to do it manually.

This issue is now resolved.

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

This issue is resolved.

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 7.2.9.

CDPD-12512: SMM cannot connect to Cloudera Manager metrics service when truststore password is empty

SMM throws the following parameter validation error when empty truststore password is used:

Truststore and password must be provided.

Now, an empty truststore password can be used in SMM without an error thrown by the validation in SMM.

CDPD-22798: Incorrect SMM producer activity

A producer in SMM might appear as active minutes after it is no longer active.

CDPD-23996: No authorization data in logs

The authorization data for users does not display in logs.

Fixed Issues in Streams Replication Manager

Review the list of Streams Replication Manager issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-22389: The SRM driver role displays as healthy, but replication fails

The SRM driver role kept running even after it encountered errors that made replication impossible. This resulted in the driver role displaying as healthy in Cloudera Manager even when no replication was done. The driver now automatically stops when such an error is encountered.

CDPD-22094: The SRM service role displays as healthy, but no metrics are processed

The SRM service role kept running even after it encountered errors that made metrics processing impossible. This resulted in the service role displaying as healthy in Cloudera Manager even when no metrics were processed. The SRM service role now automatically stops if such an error is encountered.

CDPD-23683: The replication status reported by the SRM service role for healthy replications is flaky

The replication status of healthy replications was flaky. The status of the replication could change from active to inactive frequently even if the replication was healthy. The replication status of healthy replications is no longer flaky. The correct (healthy) status is displayed when the replication is healthy.

Fixed Issues in Apache YARN

Review the list of YARN issues that are resolved in Cloudera Runtime 7.2.9.

COMPX-4550: The Hive On Tez queries fails after submitting to dynamically created pools.

When you use Hive-on-Tez with application tags, the access control check failed for dynamically created queues. With this change, the ACL settings of the parent are looked up. This issue is now resolved.

COMPX-4688: The fs2cs tool displays an exception if the weight of single leaf queue is 0 or the sum of weights is zero.

Now the fs2cs tool generates the correct capacity value in the output. This issue is now resolved.

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

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

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

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

Apache patch information

- YARN-10714
- YARN-10674

Fixed Issues in Zeppelin

Review the list of Zeppelin issues that are resolved in Cloudera Runtime 7.2.9.

CDPD-17187: Upgrade to Angular 1.8.0 due to CVEs.

Mitigate several CVE(s) including CVE-2020-7676 by upgrading to angular-1.8.

CDPD-17186: Upgrade to bootstrap 3.4.1 or 4.3.1+.

Mitigate several CVE(s) CVE-2018-14040, CVE-2018-14041, CVE-2018-14042, and CVE-2019-8331 by upgrading to bootstrap 3.4.1.

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

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

CDPD-1683: Zeppelin demo users have been removed

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

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

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

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

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

Apache patch information

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

- ZEP-97: [ZEPPELIN-3690] display all column with the same name in ui-grid
- ZEP-79: Disable fs.file.impl cache to ensure RawLocalFS is used

Fixed Issues in Apache ZooKeeper

Review the list of ZooKeeper issues that are resolved in Cloudera Runtime 7.2.9.

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

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

Apache Patch Information

- ZOOKEEPER-3037: Add JVMPauseMonitor
- ZOOKEEPER-4220: Potential redundant connection attempts during leader election

Fixed Issues In Cloudera Runtime 7.2.9.12

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.9.12.

The following issues are resolved:

- HOTREQ-1114 - Hue does not work with medium duty DL because IDBroker config has comma separated URLs

Service Pack in Cloudera Runtime 7.2.9

You can review the list of CDP Public Cloud hotfixes rolled into Cloudera Runtime 7.2.9. 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-5067

Known Issues In Cloudera Runtime 7.2.9

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.2.9.

Known Issues in Apache Atlas

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

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@58112
bc4 failed with status 404 (Not Found) Response Body
```



```
({"errorCode":"ATLAS-404-00-007","errorMessage":"Invalid instance creation/updation parameters passed : hbase_column_family.table: mandatory attribute value missing in type hbase_column_family"})
```

None

CDPD-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

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 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:

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

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

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

Workaround: N/A

Apache Issue: [HBASE-6992](#)

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

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

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

Apache Issue: N/A

Technical Service Bulletins

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

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

Upstream JIRA

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

Knowledge article

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

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.

CDPD-26975/HADOOP-17631: Using the S3A connector 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 fs.s3a.buffer.dir to \${hadoop.tmp.dir} in the site configuration.

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

None.

Unsupported Features

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

- ACLs for the NFS gateway ([HADOOP-11004](#))
- Aliyun Cloud Connector ([HADOOP-12756](#))
- Allow HDFS block replicas to be provided by an external storage system ([HDFS-9806](#))
- Consistent standby Serving reads ([HDFS-12943](#))
- Cost-Based RPC FairCallQueue ([HDFS-14403](#))
- HDFS Router Based Federation ([HDFS-10467](#))
- More than two NameNodes ([HDFS-6440](#))
- NameNode Federation ([HDFS-1052](#))
- NameNode Port-based Selective Encryption ([HDFS-13541](#))
- Non-Volatile Storage Class Memory (SCM) in HDFS Cache Directives ([HDFS-13762](#))
- OpenStack Swift ([HADOOP-8545](#))
- SFTP FileSystem ([HADOOP-5732](#))
- Storage policy satisfier ([HDFS-10285](#))

Known Issues in Apache Hive

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

CDPD-14361: Hive compaction on CDP Azure environment does not work in 7.2.9.

None.

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

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

CDPD-13636: Hive job fails with OutOfMemory exception in the Azure DE cluster

Set the parameter hive.optimize.sort.dynamic.partition.threshold=0. Add this parameter in Cloudera Manager (Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml)

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

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

appender.DRFA.strategy.action.PathConditions.nestedConditions.exceeds=same value as
appender.DRFA.strategy.max

HiveServer Web UI displays incorrect data

If you enabled auto-TLS for TLS encryption, the HiveServer2 Web UI does not display the correct data in the following tables: Active Sessions, Open Queries, Last Max n Closed Queries

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

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

Use one of the workarounds:

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

Technical Service Bulletins

TSB 2021-501: JOIN queries return wrong result for join keys with large size in Hive

JOIN queries return wrong results when performing joins on large size keys (larger than 255 bytes). This happens when the fast hash table join algorithm is enabled, which is enabled by default.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-501: JOIN queries return wrong result for join keys with large size in Hive](#)

TSB 2021-520: Cleaner causes data loss when processing an aborted dynamic partitioning transaction

Data loss may occur when an operation that involves dynamic partitioning is aborted in Hive. Cleaner does not know what partition contains the aborted deltas, so it goes over all partitions and removes aborted and `obsolete` deltas below the HighWatermark (highest writeid that could be cleaned up). Those `obsolete` deltas may be `active` ones. There is no easy way to identify obsolete deltas that are active because HighWatermark is defined on a table level.

Upstream JIRA

[HIVE-25502](#)

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: [TSB 2021-520: Cleaner causes data loss when processing an aborted dynamic partitioning transaction](#)

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

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

Knowledge article

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

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

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

Upstream JIRA

[HIVE-26235](#)

Knowledge article

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

Known Issues in Hue

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

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

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

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

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

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

3. Back up the hue directory:

```
cp -R hue hue_original
```

4. Change to the hue directory:

```
cd hue
```

5. Install the Wheel package using pip:

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

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

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

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

7. Install Tablib version 0.12.1 as follows:

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

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

Impala SELECT table query fails with UTF-8 codec error


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

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

Hue Importer is not supported in the Data Engineering template

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

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



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

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

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

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

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

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

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

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

For example, cdp-123456-scale.

Unsupported features

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

You can export Oozie workflows, schedules, and bundles from Hue and import them only within the same cluster if the cluster is unchanged. You can migrate bundle and coordinator jobs with their workflows only if their arguments have not changed between the old and the new cluster. For

example, hostnames, NameNode, Resource Manager names, YARN queue names, and all the other parameters defined in the workflow.xml and job.properties files.

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



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

Technical Service Bulletins

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

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

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

Knowledge article

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

Known Issues in Apache Impala

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

Impala known limitation when querying compacted tables

When the compaction process deletes the files for a table from the underlying HDFS location, the Impala service does not detect the changes as the compactions does not allocate new write ids. When the same table is queried from Impala it throws a 'File does not exist' exception that looks something like this:

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

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

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

An OUTER JOIN query could omit some expected result rows due to a constant such as FALSE in another join clause. For example:

```
explain SELECT 1 FROM alltypestiny a1
  INNER JOIN alltypesagg a2 ON a1.smallint_col = a2.year AND fals
e
  RIGHT JOIN alltypes a3 ON a1.year = a1.bigint_col;
+---+
| Explain String |
+---+
| Estimated Per-Host Requirements: Memory=1.00KB VCores=1 |
| 00:EMPTYSET |
+---+
```

IMPALA-3509: Breakpad minidumps can be very large when the thread count is high

The size of the breakpad minidump files grows linearly with the number of threads. By default, each thread adds 8 KB to the minidump size. Minidump files could consume significant disk space when the daemons have a high number of threads.

Add `-minidump_size_limit_hint_kb=size` to set a soft upper limit on the size of each minidump file. If the minidump file would exceed that limit, Impala reduces the amount of information for each thread from 8 KB to 2 KB. (Full thread information is captured for the first 20 threads, then 2 KB per thread after that.) The minidump file can still grow larger than the "hinted" size. For example, if you have 10,000 threads, the minidump file can be more than 20 MB.

IMPALA-4978: Impala requires FQDN from hostname command on Kerberized clusters

The method Impala uses to retrieve the host name while constructing the Kerberos principal is the `gethostname()` system call. This function might not always return the fully qualified domain name, depending on the network configuration. If the daemons cannot determine the FQDN, Impala does not start on a Kerberized cluster.

Test if a host is affected by checking whether the output of the `hostname` command includes the FQDN. On hosts where `hostname` only returns the short name, pass the command-line flag `##hostname=FULLY_QUALIFIED_DOMAIN_NAME` in the startup options of all Impala-related daemons.

IMPALA-6671: Metadata operations block read-only operations on unrelated tables

Metadata operations that change the state of a table, like `COMPUTE STATS` or `ALTER RECOVER PARTITIONS`, may delay metadata propagation of unrelated unloaded tables triggered by statements like `DESCRIBE` or `SELECT` queries.

Workaround: None

IMPALA-7072: Impala does not support Heimdal Kerberos**CDPD-28139: Set `spark.hadoop.hive.stats.autogather` to false by default**

As an Impala user, if you submit a query against a table containing data ingested using Spark and you are concerned about the quality of the query plan, you must run `COMPUTE STATS` against such a table in any case after an ETL operation because `numRows` created by Spark could be incorrect. Also, use other stats computed by `COMPUTE STATS`, e.g., Number of Distinct Values (NDV) and NULL count for good selectivity estimates.

For example, when a user ingests data from a file into a partition of an existing table using Spark, if `spark.hadoop.hive.stats.autogather` is not set to false explicitly, `numRows` associated with this partition would be 0 even though there is at least one row in the file. To avoid this, the workaround is to set `"spark.hadoop.hive.stats.autogather=false"` in the "Spark Client Advanced Configuration Snippet (Safety Valve) for `spark-conf/spark-defaults.conf`" in Spark's CM Configuration section.

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 Kafka, the impact or changes to the functionality, and the workaround.

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

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

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

Cloudera Manager > SMM > Configuration > Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml > Save Changes > Restart SMM.

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

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

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

Certain Kafka command line tools require direct access to Zookeeper

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

- `kafka-reassign-partitions`

None

The offsets.topic.replication.factor property must be less than or equal to the number of live brokers

The `offsets.topic.replication.factor` broker configuration is now enforced upon auto topic creation. Internal auto topic creation will fail with a `GROUP_COORDINATOR_NOT_AVAILABLE` error until the cluster size meets this replication factor requirement.

None

Requests fail when sending to a nonexistent topic with `auto.create.topics.enable` set to true

The first few produce requests fail when sending to a nonexistent topic with `auto.create.topics.enable` set to `true`.

Increase the number of retries in the producer configuration setting `retries`.

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

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

None

KAFKA-2561: Performance degradation when SSL Is enabled

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

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

OPSAPS-43236: Kafka garbage collection logs are written to the process directory

By default Kafka garbage collection logs are written to the agent process directory. Changing the default path for these log files is currently unsupported.

None

CDPD-24828: AvroConnectTranslator does not handle optional nested records properly

When using Avro records with optional nested records, Kafka Connect throws the following exception:

```
org.apache.kafka.connect.errors.DataException: Struct schemas do not match
```

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 Kerberos

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

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

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

Known Issues in Apache Knox

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

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

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

To prevent additional access to Atlas, close all browser windows and exit the browser.

Known Issues in Apache Kudu

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

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

None

Kudu HMS Sync is disabled and is not yet supported

None

Known Issues in Apache Oozie

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

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

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

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

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

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

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

Unsupported Feature

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

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

Cloudera does not support using Derby database with Oozie. You can use it for testing or debugging purposes, but Cloudera does not recommend using it in production environments. This could cause failures while upgrading from CDH to CDP.

BUG-123856: Upgrade fails while configuring Oozie server.

None

Known Issues in Apache Phoenix

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

CDPD-25475: The Phoenix-Hive connector cannot write to Apache Phoenix, and the relevant MapReduce job fails. When using the phoenix-hive connector, MapReduce jobs fails because the MapReduce classpath for Hive includes both the shaded and unshaded HBase JAR files causing an error.

None.

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

There are no known issues for Schema Registry in Cloudera Runtime 7.2.9.

Known Issues in Cloudera Search

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

Known Issues

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

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

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

Knowledge article

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

MapreduceIndexerTool performance problem in CDP

The reduce step of the MorphlineMapper task of the MapReduceIndexerTool (MRIT) can take very long to finish in CDPD. The reason of the slowness is merging norms without HDFS caching.

HDFS caching can not be enabled in the affected MRIT versions. Future MRIT releases will both allow controlling HDFS caching and will turn it on by default. For existing MRIT releases, the only known workaround is omitting norms. This disables length normalization for the field, saves

some memory and improves MRIT execution times. Only full-text fields can benefit from norms. Norms are omitted for primitive (non-analyzed) types by default. (Norms were formerly also used for index-time boosting but this usage has been deprecated. Index-time boosting can be achieved using doc values fields instead.)

The downside of omitting norms is that document length will not play a role in result ranking. (With norms enabled, documents with a shorter matching field would be ranked higher than matching documents with a longer field.)

You can control norms in the schema using the `omitNorms` attribute in the `fieldType` elements. To eliminate the slowdown, you must add `omitNorms="true"` to all `fieldType` elements. It is also possible to selectively set this attribute on selected fields, which allows reducing the slowdown without completely eliminating it.

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 Solr

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

Technical Service Bulletins

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

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

Knowledge article

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

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

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

Knowledge article

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

Known Issues in Apache Spark

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

CDPD-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.vectorized.use.checked.expressions does not exist
19/08/09 11:48:04 WARN conf.HiveConf: HiveConf of name hive.tez.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-3293: Cannot create views (CREATE VIEW statement) from Spark

Apache Ranger in CDP disallows Spark users from running CREATE VIEW statements.

Create the view using Hive or Impala.

CDPD-3783: Cannot create databases from Spark

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

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

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

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

Known issues in Streams Messaging Manager

Learn about the known issues for Streams Messaging Manager in Cloudera Runtime 7.2.9.

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

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

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

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

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

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

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

Workaround: 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 the SMM log:

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

Workaround: Additional security properties must be set for SMM.

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

Add the following JVM SSL properties:

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

Known Issues in Streams Replication Manager

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

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

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

None

CDPD-14019: SRM may automatically re-create deleted topics

If auto.create.topics.enable is enabled, deleted topics are automatically recreated on source clusters.

Prior to deletion, remove the topic from the topic allowlist with the srm-control tool. This prevents topics from being re-created.

```
srm-control topics --source [SOURCE_CLUSTER] --target [TARGET_CLUSTER] --remove [TOPIC1][TOPIC2]
```

CDPD-13864 and CDPD-15327: Replication stops after the network configuration of a source or target cluster is changed

If the network configuration of a cluster which is taking part in a replication flow is changed, for example, port numbers are changed as a result of enabling or disabling TLS, SRM will not update its internal configuration even if SRM is reconfigured and restarted. From SRM's perspective, it is the cluster identity that has changed. SRM cannot determine whether the new identity corresponds to the same cluster or not, only the owner or administrator of that cluster can know. In this case,

SRM tries to use the last known configuration of that cluster which might not be valid, resulting in the halt of replication.

There are three workarounds for this issue. Choose one of the following:

Increase the driver rebalance timeout

Increasing the rebalance timeout to 5 minutes (300000 ms) or longer can resolve the issue. In general a 5 minute timeout should be sufficient for most deployments. However, depending on your scenario, an even longer period might be required. Increasing the rebalance timeout might lead to increased latency when the SRM drivers stop. The cluster will be slower when it rebalances the load of the removed driver.

The rebalance timeout can be configured on a per cluster (alias) basis by adding the following to the Streams Replication Manager's Replication Configs Cloudera Manager property:

```
[***ALIAS***].rebalance.timeout.ms = [***VALUE***]
```

Replace [***ALIAS***] with a cluster alias specified in Streams Replication Manager Cluster alias. Do this for all clusters that are taking part in the replication process. When correctly configured, your configuration will have a rebalance.timeout.ms entry corresponding to each cluster (alias). For example:

```
primary.rebalance.timeout.ms = 30000
secondary.rebalance.timeout.ms = 30000
tertiary.rebalance.timeout.ms = 30000
```

After the new broker configuration is applied by SRM, the rebalance timeout can be reverted back to its original value, or removed from the configuration altogether.

Decrease replication admin timeout

Decreasing the replication admin timeout to 15 seconds (15000 ms) can resolve the issue. With higher loads, this might cause WARN messages to appear in the SRM driver log.

The admin timeout can be configured on a per replication basis by adding the following to the Streams Replication Manager's Replication Configs Cloudera Manager property:

```
[***REPLICATION***].admin.timeout.ms = [***VALUE***]
```

Replace [***REPLICATION***] with a replication specified in Streams Replication Manager's Replication Configs. Do this for all affected replications. When correctly configured, your configuration will have an admin.timeout.ms entry corresponding to each affected replication. For example:

```
primary->secondary.admin.timeout.ms = 15000
secondary->primary.admin.timeout.ms = 15000
```

After the new broker configuration is applied by SRM, the admin timeout can be reverted back to its original value, or removed from the configuration altogether.

Upgrade the brokers incrementally

Instead of switching over to the new configuration, open two separate listeners on the broker. One for the old configuration, and one for the new configuration. After updating SRM's configuration and restarting SRM, the old listener can be turned off. Non-inter-broker listeners can be configured with the dynamic configuration API of Kafka, this way not every listener change has to be followed by a restart.

CDPD-11709: 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

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
```

Limitations

SRM cannot replicate Ranger authorization policies to or from Kafka clusters

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

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

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



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

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 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 MapReduce and YARN

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

Known Issues

This issue comes up when a user tries to submit an app to a non-existing queue in the format of `parent.leaf`. The app gets submitted, gets an ID and its state is stored to the state store. After that an NPE will occur and the ResourceManager stops with the application data stored in recovery. When it gets

restarted it tries to recover the apps but once it gets to the one that was submitted to the unknown queue the NPE will happen again, causing the RM to exit again.

Resolution 1:

1. A queue needs to be created somewhere in the structure (like root.parent.leaf in case of the unknown queue was parent.leaf), this way the app still won't be launched but the NPE will not occur and the RM could start
2. the recovery data for the failing app needs to be deleted. These workarounds however do not protect the customer from experiencing the NPE again.

Resolution 2:

1. The permanent solution is to change the configuration the following way: Enable the `yarn.scheduler.capacity.queue-mappings-override.enable` (in QM as Override Queue Mappings or in capacity scheduler SV)
2. Add a new mapping rule to the first place which matches all users, places the application in the queue specified at runtime and if the target queue doesn't exist it skips this rule (goes to the next one).

Resolution 3: A custom fix can be provided as YARN-10571 is a large refactor patch.

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-5240: Restarting parent queue does not restart child queues in weight mode

When a dynamic auto child creation enabled parent queue is stopped in weight mode, its static and dynamically created child queues are also stopped. However, when the dynamic auto child creation enabled parent queue is restarted, its child queues remain stopped. In addition, the dynamically created child queues cannot be restarted manually through the YARN Queue Manager UI either.

Delete the dynamic auto child creation enabled parent queue. This action also deletes all its child queues, both static and dynamically created child queues, including the stopped dynamic queues. Then recreate the parent queue, enable the dynamic auto child creation feature for it and add any required static child queues.

COMPX-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:cdw which is less secure.

COMPX-8687: Missing access check for getAppAttempts

When the Job ACL feature is enabled using Cloudera Manager (YARN Configuration Enable JOB ACL property), the `mapreduce.cluster.acls.enabled` property is not generated to all configuration files, including the `yarn-site.xml` configuration file. As a result the ResourceManager process will use the default value of this property. The default property of `mapreduce.cluster.acls.enabled` is false.

Workaround: Enable the Job ACL feature using an advanced configuration snippet:

1. In Cloudera Manager select the YARN service.
2. Click Configuration.
3. Find the YARN Service MapReduce Advanced Configuration Snippet (Safety Valve) property.
4. Click the plus icon and add the following:
 - Name: `mapreduce.cluster.acls.enabled`
 - Value: `true`
5. Click Save Changes.

Unsupported Features

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

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

Known Issues in Apache Zeppelin

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

CDPD-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.zookeeper.ClientCnxnSocketNetty -Dzookeeper.ssl.keyStore.location=<PATH TO YOUR CONFIGURED KEYSTORE> -Dzookeeper.ssl.keyStore.password=<THE PASSWORD YOU CONFIGURED FOR THE KEYSTORE> -Dzookeeper.ssl.trustStore.location=<PATH TO YOUR CONFIGURED TRUSTSTORE> -Dzookeeper.ssl.trustStore.password=<THE PASSWORD YOU CONFIGURED FOR THE TRUSTSTORE> -Dzookeeper.client.secure=true" zookeeper-client -server <YOUR.ZOOKEEPER.SERVER-1>:2182,<YOUR.ZOOKEEPER.SERVER-2>:2182,<YOUR.ZOOKEEPER.SERVER-3>:2182
```

Behavioral Changes In Cloudera Runtime 7.2.9

You can review the changes in certain features or functionalities of components that have resulted in a change in behavior from the previously released version to this version of Cloudera Runtime 7.2.9.

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 7.2.9

Certain features and functionalities have been removed or deprecated in Cloudera Runtime 7.2.9. You must review these items to understand whether you must modify your existing configuration. You can also learn about the features that will be removed or deprecated in the future release to plan for the required changes.

Terminology

Items in this section are designated as follows:

Deprecated

Technology that Cloudera is removing in a future CDP release. Marking an item as deprecated gives you time to plan for removal in a future CDP release.

Moving

Technology that Cloudera is moving from a future CDP release and is making available through an alternative Cloudera offering or subscription. Marking an item as moving gives you time to plan for removal in a future CDP release and plan for the alternative Cloudera offering or subscription for the technology.

Removed

Technology that Cloudera has removed from CDP and is no longer available or supported as of this release. Take note of technology marked as removed since it can potentially affect your upgrade plans.

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 7.2.9. You must review these changes along with the information about the features in Kudu that will be removed or deprecated in a future release.

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

Deprecation Notices for Apache Kafka

Certain features and functionality in Kafka are deprecated or removed in Cloudera Runtime 7.2.9. You must review these changes along with the information about the features in Kafka that will be removed or deprecated in a future release.

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 7.2.9. 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 getTable (TableName) has been removed.	A client program may be interrupted by NoSuchMethodError 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 coprocessors has been changed from <code>java.util.SortedSet<E></code> to <code>org.apache.hadoop.hbase.util.SortedList<E></code> .	A client program may be interrupted by <code>NoSuchFieldError</code> exception.

MasterObserver changes

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

[#] interface `MasterObserver` (14)

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

Abstract method voidpreGetTableDescriptors (ObserverContext<MasterCoproprocessorEnvironment>, List<TableName>, List<HTableDescriptor>) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpreModifyTable (ObserverContext<MasterCoproprocessorEnvironment>, TableName, HTableDescriptor) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpreRestoreSnapshot (ObserverContext<MasterCoproprocessorEnvironment>, HBaseProtos.SnapshotDescription, HTableDescriptor) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.
Abstract method voidpreSnapshot (ObserverContext<MasterCoproprocessorEnvironment>, HBaseProtos.SnapshotDescription, HTableDescriptor) has been removed from this interface.	A client program may be interrupted by NoSuchMethodErrorException.

RegionObserver interface changes

The following changes are introduced to the RegionObserver interface.

[#] interface RegionObserver (13)

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

Abstract method `void preWALRestore (ObserverContext<RegionCoprocesorEnvironment>, HRegionInfo, HLogKey, WAL Edit)` has been removed from this interface.

A client program may be interrupted by `NoSuchMethodError` exception.

WALObserver interface changes

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

[#] interface `WALObserver`

Change	Result
Abstract method <code>void postWALWrite (ObserverContext<WALCoprocesorEnvironment>, HRegionInfo, HLogKey, WAL Edit)</code> has been removed from this interface.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>boolean preWALWrite (ObserverContext<WALCoprocesorEnvironment>, HRegionInfo, HLogKey, WAL Edit)</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 <code>void dispatch (CallRunner)</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchMethodError</code> 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 <code>void</code> to <code>boolean</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.

The following abstract methods have been removed:

[#] interface `PriorityFunction` (2)

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

Server API changes

[#] class `RpcServer` (12)

Change	Result
Type of field <code>CurCall</code> has been changed from <code>java.lang.ThreadLocal<RpcServer.Call></code> to <code>java.lang.ThreadLocal<RpcCall></code> .	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Abstract method <code>int getNumOpenConnections ()</code> has been added to this class.	This class became abstract and a client program may be interrupted by <code>InstantiationError</code> exception.
Field <code>callQueueSize</code> of type <code>org.apache.hadoop.hbase.util.Counter</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.

Field <code>connectionList</code> of type <code>java.util.List<RpcServer.Connection></code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>maxIdleTime</code> of type <code>int</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>numConnections</code> of type <code>int</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>port</code> of type <code>int</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>purgeTimeout</code> of type <code>long</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>responder</code> of type <code>RpcServer.Responder</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>socketSendBufferSize</code> of type <code>int</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.
Field <code>thresholdIdleConnections</code> of type <code>int</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchFieldError</code> exception.

Following abstract methods are removed:

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

Replication and WAL changes

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

[#] `classWALKey` (8)

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

Following fields have been removed:

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

Admin Interface API changes

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

- Thrift usage of Admin API:

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

- REST usage of Admin API:

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

[#] interface Admin (9)

Following are the changes to the Admin interface:

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

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

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

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

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

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

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

Return value type has been changed from void to java.util.concurrent.Future<java.lang.Void>.
--

This method has been removed because the return type is part of the method signature. A client program may be interrupted by NoSuchMethodError exception.

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

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

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

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

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

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

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

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

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

HTableDescriptor and HColumnDescriptor changes

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

[#] class HColumnDescriptor (1)

Change	Result
Removed super-interface 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 `ClientColumnFamilyDescriptor[]`.

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

Deprecated buffer methods

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

class OperationConflictException (1)

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

class LockTimeoutException (1)

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

Filter API changes

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

[#] class Filter (2)

Result	Result
Abstract method <code>getNextKeyHint (KeyValue)</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchMethodError</code> exception.
Abstract method <code>transform (KeyValue)</code> has been removed from this class.	A client program may be interrupted by <code>NoSuchMethodError</code> 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 <code>regionLoadPB</code> has been changed from <code>protobuf.generated.ClusterStatusProtos.RegionLoad</code> to <code>shaded.protobuf.generated.ClusterStatusProtos.RegionLoad</code> .	A client program may be interrupted by <code>NoSuchFieldError</code> 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 nextRaw (List<Cell>, int) has been removed from this interface.	A client program may be interrupted by NoSuchMethodError exception.

StoreFile changes

[#] class StoreFile (1)

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

MapReduce changes

HFile*Format has been removed.

ClusterStatus changes

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

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

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

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

Purge of Protocol Buffers from API

Protocol Buffers (PB) has been deprecated in APIs.

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

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

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

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

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

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

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

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 NoSuchElementException.