Cloudera Runtime 7.2.15

Release Notes

Date published: 2022-05-12
Date modified: 2023-01-12

https://docs.cloudera.com/

Legal Notice

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.

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

Contents

Overview

You can review the Release Notes of Cloudera Runtime 7.2.15 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 7.2.15 distribution to ensure compatibility of these components with other applications. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Apache Components

Component	Version
Apache Arrow	0.11.1.7.2.15.0-147
Apache Atlas	2.1.0.7.2.15.0-147
Apache Calcite	1.21.0.7.2.15.0-147
Apache Avro	1.8.2.7.2.15.0-147
Apache Flink	1.14.0.1.6.1.0
Apache Hadoop (Includes YARN and HDFS)	3.1.1.7.2.15.0-147
Apache HBase	2.4.6.7.2.15.0-147
Apache Hive	3.1.3000.7.2.15.0-147
Apache Impala	4.0.0.7.2.15.0-147
Apache Kafka	2.8.1.7.2.15.0-147
Apache Knox	1.3.0.7.2.15.0-147
Apache Kudu	1.15.0.7.2.15.0-147
Apache Livy	0.6.0.7.2.15.0-147
Apache MapReduce	3.1.1.7.2.15.0-147
Apache NiFi	1.16.0.2.2.5.0
Apache NiFi Registry	1.16.0.2.2.5.0
Apache Oozie	5.1.0.7.2.15.0-147
Apache ORC	1.5.1.7.2.15.0-147
Apache Parquet	1.10.99.7.2.15.0-147
Apache Phoenix	5.1.1.7.2.15.0-147
Apache Ranger	2.1.0.7.2.15.0-147
Apache Solr	8.4.1.7.2.15.0-147
Apache Spark	2.4.8.7.2.15.0-147
Apache Spark 3	3.2.1.7.2.15.0-147
Apache Sqoop	1.4.7.7.2.15.0-147
Apache Tez	0.9.1.7.2.15.0-147
Apache Zeppelin	0.8.2.7.2.15.0-147

Component	Version
Apache ZooKeeper	3.5.5.7.2.15.0-147

Other Components

Component	Version
Cruise Control	2.0.100.7.1.7.0-551
Data Analytics Studio	1.4.2.7.2.15.0-147
GCS Connector	2.1.2.7.2.15.0-147
HBase Indexer	1.5.0.7.2.15.0-147
Hive Solr Connector	4.0.0.7.2.15.0-147
Hue	4.5.0.7.2.15.0-147
Search	1.0.0.7.2.15.0-147
Schema Registry	0.10.0.7.2.15.0-147
Spark Solr Connector	3.9.0.7.2.15.0-147
Streams Messaging Manager	2.2.0.7.2.15.0-147
Streams Replication Manager	1.1.0.7.2.15.0-147

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.2.15.0-147
Hive Meta Store (HMS)	1.0.0.7.2.15.0-147
Hive on Tez	1.0.0.7.2.15.0-147
Hive Warehouse Connector	1.0.0.7.2.15.0-147
Spark Atlas Connector	0.1.0.7.2.15.0-147
Spark Schema Registry	1.1.0.7.2.15.0-147

Using the Cloudera Runtime Maven repository 7.2.15

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

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

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

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

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

Maven Artifacts for Cloudera Runtime 7.2.15

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

Project	groupId	artifactId	version
Apache Atlas	org.apache.atlas	atlas-authorization	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-aws-s3-bridge	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-azure-adls-bridge	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-classification-updater	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-client-common	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-client-v1	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-client-v2	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-common	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-distro	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-docs	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-graphdb-api	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-graphdb-common	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-graphdb-janus	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-index-repair-tool	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-intg	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-janusgraph-hbase2	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-notification	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-plugin-classloader	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-repository	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-server-api	2.1.0.7.2.15.0-147
	org.apache.atlas	atlas-testtools	2.1.0.7.2.15.0-147
	org.apache.atlas	hbase-bridge	2.1.0.7.2.15.0-147
	org.apache.atlas	hbase-bridge-shim	2.1.0.7.2.15.0-147
	org.apache.atlas	hbase-testing-util	2.1.0.7.2.15.0-147
	org.apache.atlas	hdfs-model	2.1.0.7.2.15.0-147
	org.apache.atlas	hive-bridge	2.1.0.7.2.15.0-147
	org.apache.atlas	hive-bridge-shim	2.1.0.7.2.15.0-147
	org.apache.atlas	impala-bridge	2.1.0.7.2.15.0-147
	org.apache.atlas	impala-bridge-shim	2.1.0.7.2.15.0-147
	org.apache.atlas	impala-hook-api	2.1.0.7.2.15.0-147
	org.apache.atlas	kafka-bridge	2.1.0.7.2.15.0-147
	org.apache.atlas	kafka-bridge-shim	2.1.0.7.2.15.0-147
	org.apache.atlas	navigator-to-atlas	2.1.0.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.atlas	sample-app	2.1.0.7.2.15.0-147
	org.apache.atlas	sqoop-bridge	2.1.0.7.2.15.0-147
	org.apache.atlas	sqoop-bridge-shim	2.1.0.7.2.15.0-147
Apache Avro	org.apache.avro	avro	1.8.2.7.2.15.0-147
	org.apache.avro	avro-compiler	1.8.2.7.2.15.0-147
	org.apache.avro	avro-ipc	1.8.2.7.2.15.0-147
	org.apache.avro	avro-mapred	1.8.2.7.2.15.0-147
	org.apache.avro	avro-maven-plugin	1.8.2.7.2.15.0-147
	org.apache.avro	avro-protobuf	1.8.2.7.2.15.0-147
	org.apache.avro	avro-service-archetype	1.8.2.7.2.15.0-147
	org.apache.avro	avro-thrift	1.8.2.7.2.15.0-147
	org.apache.avro	avro-tools	1.8.2.7.2.15.0-147
	org.apache.avro	trevni-avro	1.8.2.7.2.15.0-147
	org.apache.avro	trevni-core	1.8.2.7.2.15.0-147
Apache Calcite	org.apache.calcite	calcite-babel	1.21.0.7.2.15.0-147
	org.apache.calcite	calcite-core	1.21.0.7.2.15.0-147
	org.apache.calcite	calcite-druid	1.21.0.7.2.15.0-147
	org.apache.calcite	calcite-kafka	1.21.0.7.2.15.0-147
	org.apache.calcite	calcite-linq4j	1.21.0.7.2.15.0-147
	org.apache.calcite	calcite-server	1.21.0.7.2.15.0-147
	org.apache.calcite	calcite-ubenchmark	1.21.0.7.2.15.0-147
	org.apache.calcite.avatica	avatica	1.16.0.7.2.15.0-147
	org.apache.calcite.avatica	avatica-core	1.16.0.7.2.15.0-147
	org.apache.calcite.avatica	avatica-metrics	1.16.0.7.2.15.0-147
	org.apache.calcite.avatica	avatica-metrics-dropwizardmetrics	1.16.0.7.2.15.0-147
	org.apache.calcite.avatica	avatica-noop-driver	1.16.0.7.2.15.0-147
	org.apache.calcite.avatica	avatica-server	1.16.0.7.2.15.0-147
	org.apache.calcite.avatica	avatica-standalone-server	1.16.0.7.2.15.0-147
	org.apache.calcite.avatica	avatica-tck	1.16.0.7.2.15.0-147
Apache Druid	org.apache.druid	druid-aws-common	0.17.1.7.2.15.0-147
	org.apache.druid	druid-benchmarks	0.17.1.7.2.15.0-147
	org.apache.druid	druid-console	0.17.1.7.2.15.0-147
	org.apache.druid	druid-core	0.17.1.7.2.15.0-147
	org.apache.druid	druid-gcp-common	0.17.1.7.2.15.0-147
	org.apache.druid	druid-hll	0.17.1.7.2.15.0-147
	org.apache.druid	druid-indexing-hadoop	0.17.1.7.2.15.0-147
	org.apache.druid	druid-indexing-service	0.17.1.7.2.15.0-147
	org.apache.druid	druid-integration-tests	0.17.1.7.2.15.0-147
	org.apache.druid	druid-processing	0.17.1.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.druid	druid-server	0.17.1.7.2.15.0-147
	org.apache.druid	druid-services	0.17.1.7.2.15.0-147
	org.apache.druid	druid-sql	0.17.1.7.2.15.0-147
	org.apache.druid	extendedset	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-avro-extensions	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-basic-security	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-bloom-filter	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-datasketches	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-ec2-extensions	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-google-extensions	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-hdfs-storage	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-histogram	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-kafka-extraction-namespace	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-kafka-indexing-service	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-kerberos	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-kinesis-indexing-service	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-lookups-cached-global	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-lookups-cached-single	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-orc-extensions	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-parquet-extensions	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-protobuf-extensions	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-s3-extensions	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	druid-stats	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	mysql-metadata-storage	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	postgresql-metadata-storage	0.17.1.7.2.15.0-147
	org.apache.druid.extensions	simple-client-sslcontext	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	ambari-metrics-emitter	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	dropwizard-emitter	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	druid-azure-extensions	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	druid-cassandra-storage	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	druid-cloudfiles-extensions	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	druid-distinctcount	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	druid-influx-extensions	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	druid-influxdb-emitter	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	druid-momentsketch	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	druid-moving-average-query	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	druid-opentsdb-emitter	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	druid-redis-cache	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	druid-tdigestsketch	0.17.1.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.druid.extensions.contrib	druid-thrift-extensions	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	druid-time-min-max	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	druid-virtual-columns	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	graphite-emitter	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	kafka-emitter	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	materialized-view-maintenance	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	materialized-view-selection	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	sqlserver-metadata-storage	0.17.1.7.2.15.0-147
	org.apache.druid.extensions.contrib	statsd-emitter	0.17.1.7.2.15.0-147
GCS Connector	com.google.cloud.bigdataoss	bigquery-connector	2.1.2.7.2.15.0-147
	com.google.cloud.bigdataoss	gcs-connector	2.1.2.7.2.15.0-147
	com.google.cloud.bigdataoss	gcsio	2.1.2.7.2.15.0-147
	com.google.cloud.bigdataoss	util	2.1.2.7.2.15.0-147
	com.google.cloud.bigdataoss	util-hadoop	2.1.2.7.2.15.0-147
Apache Hadoop	org.apache.hadoop	hadoop-aliyun	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-annotations	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-archive-logs	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-archives	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-assemblies	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-auth	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-aws	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-azure	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-azure-datalake	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-build-tools	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-client	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-client-api	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-client-integration-tests	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-client-minicluster	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-client-runtime	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-cloud-storage	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-common	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-datajoin	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-distcp	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-extras	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-fs2img	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-gridmix	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-hdds-client	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-hdds-common	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-hdds-config	1.1.0.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-hdds-container-service	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-hdds-docs	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-hdds-hadoop-dependency-client	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-hdds-hadoop-dependency-server	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-hdds-hadoop-dependency-test	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-hdds-interface-admin	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-hdds-interface-client	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-hdds-interface-server	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-hdds-server-framework	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-hdds-server-scm	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-hdds-test-utils	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-hdds-tools	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-hdfs	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-hdfs-client	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-hdfs-httpfs	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-hdfs-native-client	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-hdfs-nfs	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-hdfs-rbf	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-kafka	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-kms	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-mapreduce-client-app	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-mapreduce-client-common	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-mapreduce-client-core	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-mapreduce-client-hs	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-mapreduce-examples	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-maven-plugins	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-minicluster	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-minikdc	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-nfs	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-openstack	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-client	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-common	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-csi	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-datanode	1.1.0.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-ozone-dist	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-filesystem	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-filesystem-common	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-filesystem-hadoop2	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-filesystem-hadoop3	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-filesystem-shaded	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-insight	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-integration-test	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-interface-client	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-interface-storage	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-network-tests	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-ozone-manager	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-recon	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-reconcodegen	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-s3gateway	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-ozone-tools	1.1.0.7.2.15.0-147
	org.apache.hadoop	hadoop-resourceestimator	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-rumen	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-sls	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-streaming	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-tools-dist	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-api	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-client	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-common	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-registry	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-server-common	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-server-router	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-server-tests	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.1.1.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-services-api	3.1.1.7.2.15.0-147
	org.apache.hadoop	hadoop-yarn-services-core	3.1.1.7.2.15.0-147
	org.apache.hadoop	mini-chaos-tests	1.1.0.7.2.15.0-147
Apache HBase	org.apache.hbase	filesystem	hadoop3-3-testutils
	org.apache.hbase	hbase-annotations	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-asyncfs	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-checkstyle	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-client	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-client-project	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-common	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-endpoint	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-examples	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-external-blockcache	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-hadoop-compat	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-hadoop2-compat	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-hbtop	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-http	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-it	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-logging	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-mapreduce	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-metrics	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-metrics-api	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-procedure	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-protocol	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-protocol-shaded	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-replication	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-resource-bundle	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-rest	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-rsgroup	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-server	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-shaded-client	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-shaded-client-project	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-shaded-mapreduce	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-shaded-testing-util	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-shaded-testing-util-tester	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-shell	2.4.6.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.hbase	hbase-testing-util	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-thrift	2.4.6.7.2.15.0-147
	org.apache.hbase	hbase-zookeeper	2.4.6.7.2.15.0-147
	org.apache.hbase.connectors.kafka	hbase-kafka-model	1.0.0.7.2.15.0-147
	org.apache.hbase.connectors.kafka	hbase-kafka-proxy	1.0.0.7.2.15.0-147
	org.apache.hbase.connectors.spark	hbase-spark	1.0.0.7.2.15.0-147
	org.apache.hbase.connectors.spark	hbase-spark-it	1.0.0.7.2.15.0-147
	org.apache.hbase.connectors.spark	hbase-spark-protocol	1.0.0.7.2.15.0-147
	org.apache.hbase.connectors.spark	hbase-spark-protocol-shaded	1.0.0.7.2.15.0-147
	org.apache.hbase.connectors.spark	hbase-spark3	1.0.0.7.2.15.0-147
	org.apache.hbase.connectors.spark	hbase-spark3-it	1.0.0.7.2.15.0-147
	org.apache.hbase.connectors.spark	hbase-spark3-protocol	1.0.0.7.2.15.0-147
	org.apache.hbase.connectors.spark	hbase-spark3-protocol-shaded	1.0.0.7.2.15.0-147
	org.apache.hbase.filesystem	hadoop-testutils	1.0.0.7.2.15.0-147
	org.apache.hbase.filesystem	hbase-fs-impl	1.0.0.7.2.15.0-147
	org.apache.hbase.filesystem	hbase-oss	1.0.0.7.2.15.0-147
	org.apache.hbase.thirdparty	hbase-noop-htrace	3.5.0.7.2.15.0-147
	org.apache.hbase.thirdparty	hbase-shaded-gson	3.5.0.7.2.15.0-147
	org.apache.hbase.thirdparty	hbase-shaded-jersey	3.5.0.7.2.15.0-147
	org.apache.hbase.thirdparty	hbase-shaded-jetty	3.5.0.7.2.15.0-147
	org.apache.hbase.thirdparty	hbase-shaded-miscellaneous	3.5.0.7.2.15.0-147
	org.apache.hbase.thirdparty	hbase-shaded-netty	3.5.0.7.2.15.0-147
	org.apache.hbase.thirdparty	hbase-shaded-protobuf	3.5.0.7.2.15.0-147
Apache Hive	org.apache.hive	catalogd-unit	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-beeline	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-blobstore	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-classification	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-cli	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-common	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-contrib	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-druid-handler	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-exec	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-hbase-handler	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-hplsql	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-impala	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-it-custom-serde	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-it-druid	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-it-impala	3.1.3000.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.hive	hive-it-minikdc	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-it-qfile	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-it-test-serde	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-it-unit	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-it-util	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-jdbc	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-jdbc-handler	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-jmh	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-kryo-registrator	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-kudu-handler	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-llap-client	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-llap-common	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-llap-ext-client	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-llap-server	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-llap-tez	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-metastore	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-parser	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-pre-upgrade	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-serde	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-service	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-service-rpc	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-shims	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-spark-client	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-standalone-metastore	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-storage-api	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-streaming	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-testutils	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-udf	3.1.3000.7.2.15.0-147
	org.apache.hive	hive-vector-code-gen	3.1.3000.7.2.15.0-147
	org.apache.hive	kafka-handler	3.1.3000.7.2.15.0-147
	org.apache.hive.hcatalog	hive-hcatalog-core	3.1.3000.7.2.15.0-147
	org.apache.hive.hcatalog	hive-hcatalog-pig-adapter	3.1.3000.7.2.15.0-147
	org.apache.hive.hcatalog	hive-hcatalog-server-extensions	3.1.3000.7.2.15.0-147
	org.apache.hive.hcatalog	hive-hcatalog-streaming	3.1.3000.7.2.15.0-147
	org.apache.hive.hcatalog	hive-webhcat	3.1.3000.7.2.15.0-147
	org.apache.hive.hcatalog	hive-webhcat-java-client	3.1.3000.7.2.15.0-147
	org.apache.hive.hive-it-custom-udfs	udf-classloader-udf1	3.1.3000.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.hive.hive-it-custom-udfs	udf-classloader-udf2	3.1.3000.7.2.15.0-147
	org.apache.hive.hive-it-custom-udfs	udf-classloader-util	3.1.3000.7.2.15.0-147
	org.apache.hive.hive-it-custom-udfs	udf-vectorized-badexample	3.1.3000.7.2.15.0-147
	org.apache.hive.shims	hive-shims-0.23	3.1.3000.7.2.15.0-147
	org.apache.hive.shims	hive-shims-common	3.1.3000.7.2.15.0-147
	org.apache.hive.shims	hive-shims-scheduler	3.1.3000.7.2.15.0-147
Apache Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector_2.11	1.0.0.7.2.15.0-147
Apache Kafka	org.apache.kafka	connect	2.8.1.7.2.15.0-147
	org.apache.kafka	connect-api	2.8.1.7.2.15.0-147
	org.apache.kafka	connect-basic-auth-extension	2.8.1.7.2.15.0-147
	org.apache.kafka	connect-cloudera-authorization-extension	2.8.1.7.2.15.0-147
	org.apache.kafka	connect-cloudera-common	2.8.1.7.2.15.0-147
	org.apache.kafka	connect-cloudera-secret-storage	2.8.1.7.2.15.0-147
	org.apache.kafka	connect-cloudera-security-policies	2.8.1.7.2.15.0-147
	org.apache.kafka	connect-file	2.8.1.7.2.15.0-147
	org.apache.kafka	connect-json	2.8.1.7.2.15.0-147
	org.apache.kafka	connect-mirror	2.8.1.7.2.15.0-147
	org.apache.kafka	connect-mirror-client	2.8.1.7.2.15.0-147
	org.apache.kafka	connect-runtime	2.8.1.7.2.15.0-147
	org.apache.kafka	connect-transforms	2.8.1.7.2.15.0-147
	org.apache.kafka	generator	2.8.1.7.2.15.0-147
	org.apache.kafka	jmh-benchmarks	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-clients	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-cloudera-plugins	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-examples	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-log4j-appender	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-metadata	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-raft	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-shell	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-examples	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-scala_2.12	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-scala_2.13	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-test-utils	2.8.1.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-upgrade-system-tests-10	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-upgrade-system-tests-11	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-upgrade-system-tests-20	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-upgrade-system-tests-21	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-upgrade-system-tests-22	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-upgrade-system-tests-23	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-upgrade-system-tests-24	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-upgrade-system-tests-25	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-upgrade-system-tests-26	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-streams-upgrade-system-tests-27	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka-tools	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka_2.12	2.8.1.7.2.15.0-147
	org.apache.kafka	kafka_2.13	2.8.1.7.2.15.0-147
	org.apache.kafka	ranger-kafka-connect-plugin	2.8.1.7.2.15.0-147
Apache Knox	org.apache.knox	gateway-adapter	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-admin-ui	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-applications	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-cloud-bindings	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-demo-ldap	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-demo-ldap-launcher	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-discovery-ambari	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-discovery-cm	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-docker	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-i18n	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-i18n-logging-log4j	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-i18n-logging-sl4j	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-performance-test	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-ha	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-identity-assertion-common	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-identity-assertion-concat	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-identity-assertion-no-doas	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-identity-assertion-pseudo	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-identity-assertion-regex	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-identity-assertion-switchcase	1.3.0.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.knox	gateway-provider-jersey	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-rewrite	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-rewrite-common	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-rewrite-func-service-registry	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-rewrite-step-secure-query	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-security-authc-anon	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-security-authz-acls	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-security-authz-composite	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-security-clientcert	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-security-hadoopauth	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-security-jwt	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-security-pac4j	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-security-preauth	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-security-shiro	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-provider-security-webappsec	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-release	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-server	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-server-launcher	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-server-xforwarded-filter	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-admin	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-as	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-definitions	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-hashicorp-vault	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-hbase	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-health	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-hive	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-idbroker	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-impala	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-jkg	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-knoxsso	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-knoxssout	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-knoxtoken	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-livy	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-metadata	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-nifi	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-nifi-registry	1.3.0.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.knox	gateway-service-remoteconfig	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-rm	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-session	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-storm	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-test	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-tgs	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-vault	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-service-webhdfs	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-shell	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-shell-launcher	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-shell-release	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-shell-samples	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-spi	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-test	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-test-idbroker	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-test-release-utils	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-test-utils	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-topology-hadoop-xml	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-topology-simple	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-util-common	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-util-configinjector	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-util-launcher	1.3.0.7.2.15.0-147
	org.apache.knox	gateway-util-urltemplate	1.3.0.7.2.15.0-147
	org.apache.knox	hadoop-examples	1.3.0.7.2.15.0-147
	org.apache.knox	knox-cli-launcher	1.3.0.7.2.15.0-147
	org.apache.knox	knox-homepage-ui	1.3.0.7.2.15.0-147
	org.apache.knox	knox-token-management-ui	1.3.0.7.2.15.0-147
	org.apache.knox	webhdfs-kerb-test	1.3.0.7.2.15.0-147
	org.apache.knox	webhdfs-test	1.3.0.7.2.15.0-147
Apache Kudu	org.apache.kudu	kudu-backup-tools	1.15.0.7.2.15.0-147
	org.apache.kudu	kudu-backup2_2.11	1.15.0.7.2.15.0-147
	org.apache.kudu	kudu-backup3_2.12	1.15.0.7.2.15.0-147
	org.apache.kudu	kudu-client	1.15.0.7.2.15.0-147
	org.apache.kudu	kudu-hive	1.15.0.7.2.15.0-147
	org.apache.kudu	kudu-spark2-tools_2.11	1.15.0.7.2.15.0-147
	org.apache.kudu	kudu-spark2_2.11	1.15.0.7.2.15.0-147
	org.apache.kudu	kudu-spark3-tools_2.12	1.15.0.7.2.15.0-147
	org.apache.kudu	kudu-spark3_2.12	1.15.0.7.2.15.0-147
	org.apache.kudu	kudu-test-utils	1.15.0.7.2.15.0-147

Project	groupId	artifactId	version
Apache Livy	org.apache.livy	livy-api	0.6.0.7.2.15.0-147
	org.apache.livy	livy-api	0.6.3000.7.2.15.0-147
	org.apache.livy	livy-client-common	0.6.0.7.2.15.0-147
	org.apache.livy	livy-client-http	0.6.3000.7.2.15.0-147
	org.apache.livy	livy-core_2.11	0.6.0.7.2.15.0-147
	org.apache.livy	livy-core_2.12	0.6.3000.7.2.15.0-147
	org.apache.livy	livy-examples	0.6.3000.7.2.15.0-147
	org.apache.livy	livy-integration-test	0.6.3000.7.2.15.0-147
	org.apache.livy	livy-repl_2.11	0.6.3000.7.2.15.0-147
	org.apache.livy	livy-repl_2.12	0.6.3000.7.2.15.0-147
	org.apache.livy	livy-rsc	0.6.3000.7.2.15.0-147
	org.apache.livy	livy-scala-api_2.11	0.6.0.7.2.15.0-147
	org.apache.livy	livy-scala-api_2.12	0.6.3000.7.2.15.0-147
	org.apache.livy	livy-server	0.6.0.7.2.15.0-147
	org.apache.livy	livy-test-lib	0.6.0.7.2.15.0-147
	org.apache.livy	livy-thriftserver	0.6.0.7.2.15.0-147
	org.apache.livy	livy-thriftserver-session	0.6.0.7.2.15.0-147
Apache Lucene	org.apache.lucene	lucene-analyzers-common	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-analyzers-icu	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-analyzers-kuromoji	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-analyzers-morfologik	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-analyzers-nori	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-analyzers-opennlp	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-analyzers-phonetic	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-analyzers-smartcn	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-analyzers-stempel	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-backward-codecs	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-benchmark	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-classification	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-codecs	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-core	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-demo	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-expressions	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-facet	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-grouping	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-highlighter	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-join	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-memory	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-misc	8.4.1.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.lucene	lucene-monitor	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-queries	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-queryparser	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-replicator	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-sandbox	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-spatial	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-spatial-extras	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-spatial3d	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-suggest	8.4.1.7.2.15.0-147
	org.apache.lucene	lucene-test-framework	8.4.1.7.2.15.0-147
Apache Oozie	org.apache.oozie	oozie-client	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-core	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-distro	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-examples	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-server	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-sharelib-git	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-sharelib-spark	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-tools	5.1.0.7.2.15.0-147
	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.2.15.0-147
	org.apache.oozie.test	oozie-mini	5.1.0.7.2.15.0-147
Apache ORC	org.apache.orc	orc-core	1.5.1.7.2.15.0-147
	org.apache.orc	orc-examples	1.5.1.7.2.15.0-147
	org.apache.orc	orc-mapreduce	1.5.1.7.2.15.0-147
	org.apache.orc	orc-shims	1.5.1.7.2.15.0-147
	org.apache.orc	orc-tools	1.5.1.7.2.15.0-147
Apache Parquet	org.apache.parquet	parquet-avro	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-cascading	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-cascading3	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-column	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-common	1.10.99.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.parquet	parquet-encoding	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-format-structures	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-generator	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-hadoop	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-hadoop-bundle	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-jackson	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-pig	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-pig-bundle	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-protobuf	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-scala_2.10	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-thrift	1.10.99.7.2.15.0-147
	org.apache.parquet	parquet-tools	1.10.99.7.2.15.0-147
Apache Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.4	5.1.1.7.2.15.0-147
	org.apache.phoenix	phoenix-client-hbase-2.4	5.1.1.7.2.15.0-147
	org.apache.phoenix	phoenix-connectors-phoenix5-compat	6.0.0.7.2.15.0-147
	org.apache.phoenix	phoenix-core	5.1.1.7.2.15.0-147
	org.apache.phoenix	phoenix-hbase-compat-2.1.6	5.1.1.7.2.15.0-147
	org.apache.phoenix	phoenix-hbase-compat-2.2.5	5.1.1.7.2.15.0-147
	org.apache.phoenix	phoenix-hbase-compat-2.3.0	5.1.1.7.2.15.0-147
	org.apache.phoenix	phoenix-hbase-compat-2.4.0	5.1.1.7.2.15.0-147
	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.1.1.7.2.15.0-147
	org.apache.phoenix	phoenix-pherf	5.1.1.7.2.15.0-147
	org.apache.phoenix	phoenix-queryserver	6.0.0.7.2.15.0-147
	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.2.15.0-147
	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.2.15.0-147
	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.2.15.0-147
	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.2.15.0-147
	org.apache.phoenix	phoenix-server-hbase-2.4	5.1.1.7.2.15.0-147
	org.apache.phoenix	phoenix-tracing-webapp	5.1.1.7.2.15.0-147
	org.apache.phoenix	phoenix5-hive	6.0.0.7.2.15.0-147
	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.2.15.0-147
	org.apache.phoenix	phoenix5-spark	6.0.0.7.2.15.0-147
	org.apache.phoenix	phoenix5-spark-shaded	6.0.0.7.2.15.0-147
	org.apache.phoenix	phoenix5-spark3	6.0.0.7.2.15.0-147
	org.apache.phoenix	phoenix5-spark3-shaded	6.0.0.7.2.15.0-147
	org.apache.phoenix.thirdparty	phoenix-shaded-commons-cli	1.1.0.7.2.15.0-147
	org.apache.phoenix.thirdparty	phoenix-shaded-guava	1.1.0.7.2.15.0-147
Apache Ranger	org.apache.ranger	conditions-enrichers	2.1.0.7.2.15.0-147
	org.apache.ranger	credentialbuilder	2.1.0.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.ranger	embeddedwebserver	2.1.0.7.2.15.0-147
	org.apache.ranger	jisql	2.1.0.7.2.15.0-147
	org.apache.ranger	ldapconfigcheck	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-adls-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-atlas-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-atlas-plugin-shim	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-authn	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-distro	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-examples-distro	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-hbase-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-hbase-plugin-shim	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-hdfs-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-hdfs-plugin-shim	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-hive-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-hive-plugin-shim	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-intg	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-kafka-connect-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-kafka-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-kafka-plugin-shim	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-kms	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-kms-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-kms-plugin-shim	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-knox-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-knox-plugin-shim	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-kudu-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-kylin-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-kylin-plugin-shim	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-nifi-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-nifi-registry-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-ozone-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-ozone-plugin-shim	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-plugin-classloader	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-plugins-audit	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-plugins-common	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-plugins-cred	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-plugins-installer	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-policymigration	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-raz-adls	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-raz-chained-plugins	2.1.0.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.ranger	ranger-raz-hook-abfs	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-raz-hook-s3	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-raz-intg	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-raz-processor	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-raz-s3	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-raz-s3-lib	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-rms-common	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-rms-hive	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-rms-plugins-common	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-rms-webapp	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-s3-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-sampleapp-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-schema-registry-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-solr-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-solr-plugin-shim	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-sqoop-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-sqoop-plugin-shim	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-storm-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-storm-plugin-shim	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-tagsync	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-tools	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-util	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-yarn-plugin	2.1.0.7.2.15.0-147
	org.apache.ranger	ranger-yarn-plugin-shim	2.1.0.7.2.15.0-147
	org.apache.ranger	sample-client	2.1.0.7.2.15.0-147
	org.apache.ranger	sampleapp	2.1.0.7.2.15.0-147
	org.apache.ranger	shaded-raz-hook-abfs	2.1.0.7.2.15.0-147
	org.apache.ranger	shaded-raz-hook-s3	2.1.0.7.2.15.0-147
	org.apache.ranger	ugsync-util	2.1.0.7.2.15.0-147
	org.apache.ranger	unixauthclient	2.1.0.7.2.15.0-147
	org.apache.ranger	unixauthservice	2.1.0.7.2.15.0-147
	org.apache.ranger	unixusersync	2.1.0.7.2.15.0-147
Apache Solr	org.apache.solr	solr-analysis-extras	8.4.1.7.2.15.0-147
	org.apache.solr	solr-analytics	8.4.1.7.2.15.0-147
	org.apache.solr	solr-cell	8.4.1.7.2.15.0-147
	org.apache.solr	solr-clustering	8.4.1.7.2.15.0-147
	org.apache.solr	solr-core	8.4.1.7.2.15.0-147
	org.apache.solr	solr-dataimporthandler	8.4.1.7.2.15.0-147
	org.apache.solr	solr-dataimporthandler-extras	8.4.1.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.solr	solr-jaegertracer-configurator	8.4.1.7.2.15.0-147
	org.apache.solr	solr-langid	8.4.1.7.2.15.0-147
	org.apache.solr	solr-ltr	8.4.1.7.2.15.0-147
	org.apache.solr	solr-prometheus-exporter	8.4.1.7.2.15.0-147
	org.apache.solr	solr-security-util	8.4.1.7.2.15.0-147
	org.apache.solr	solr-solrj	8.4.1.7.2.15.0-147
	org.apache.solr	solr-test-framework	8.4.1.7.2.15.0-147
	org.apache.solr	solr-velocity	8.4.1.7.2.15.0-147
Apache Spark	org.apache.spark	spark-avro_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-avro_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-catalyst_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-catalyst_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-core_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-core_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-graphx_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-graphx_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-hadoop-cloud_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-hadoop-cloud_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-hive_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-hive_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-kubernetes_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-kubernetes_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-kvstore_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-kvstore_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-launcher_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-launcher_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-mllib-local_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-mllib-local_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-mllib_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-mllib_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-network-common_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-network-common_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-network-shuffle_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-network-shuffle_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-network-yarn_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-network-yarn_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-repl_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-repl_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-shaded-raz	3.2.1.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.spark	spark-sketch_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-sketch_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-sql-kafka-0-10_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-sql-kafka-0-10_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-sql_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-sql_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-streaming-kafka-0-10_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-streaming-kafka-0-10_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-streaming_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-streaming_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-tags_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-tags_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-token-provider-kafka-0-10_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-token-provider-kafka-0-10_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-unsafe_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-unsafe_2.12	3.2.1.7.2.15.0-147
	org.apache.spark	spark-yarn_2.11	2.4.8.7.2.15.0-147
	org.apache.spark	spark-yarn_2.12	3.2.1.7.2.15.0-147
Apache Sqoop	org.apache.sqoop	sqoop	1.4.7.7.2.15.0-147
	org.apache.sqoop	sqoop-test	1.4.7.7.2.15.0-147
Apache Tez	org.apache.tez	hadoop-shim	0.9.1.7.2.15.0-147
	org.apache.tez	hadoop-shim-2.8	0.9.1.7.2.15.0-147
	org.apache.tez	tez-api	0.9.1.7.2.15.0-147
	org.apache.tez	tez-aux-services	0.9.1.7.2.15.0-147
	org.apache.tez	tez-common	0.9.1.7.2.15.0-147
	org.apache.tez	tez-dag	0.9.1.7.2.15.0-147
	org.apache.tez	tez-examples	0.9.1.7.2.15.0-147
	org.apache.tez	tez-ext-service-tests	0.9.1.7.2.15.0-147
	org.apache.tez	tez-history-parser	0.9.1.7.2.15.0-147
	org.apache.tez	tez-javadoc-tools	0.9.1.7.2.15.0-147
	org.apache.tez	tez-job-analyzer	0.9.1.7.2.15.0-147
	org.apache.tez	tez-mapreduce	0.9.1.7.2.15.0-147
	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.2.15.0-147
	org.apache.tez	tez-runtime-internals	0.9.1.7.2.15.0-147
	org.apache.tez	tez-runtime-library	0.9.1.7.2.15.0-147
	org.apache.tez	tez-tests	0.9.1.7.2.15.0-147
	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.2.15.0-147

Project	groupId	artifactId	version
	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.2.15.0-147
	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.2.15.0-147
	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.2.15.0-147
Apache Zeppelin	org.apache.zeppelin	zeppelin-angular	0.8.2.7.2.15.0-147
	org.apache.zeppelin	zeppelin-display	0.8.2.7.2.15.0-147
	org.apache.zeppelin	zeppelin-interpreter	0.8.2.7.2.15.0-147
	org.apache.zeppelin	zeppelin-jdbc	0.8.2.7.2.15.0-147
	org.apache.zeppelin	zeppelin-jupyter	0.8.2.7.2.15.0-147
	org.apache.zeppelin	zeppelin-livy	0.8.2.7.2.15.0-147
	org.apache.zeppelin	zeppelin-markdown	0.8.2.7.2.15.0-147
	org.apache.zeppelin	zeppelin-server	0.8.2.7.2.15.0-147
	org.apache.zeppelin	zeppelin-shaded-raz	0.8.2.7.2.15.0-147
	org.apache.zeppelin	zeppelin-shell	0.8.2.7.2.15.0-147
	org.apache.zeppelin	zeppelin-zengine	0.8.2.7.2.15.0-147
Apache ZooKeeper	org.apache.zookeeper	zookeeper	3.5.5.7.2.15.0-147
	org.apache.zookeeper	zookeeper-client-c	3.5.5.7.2.15.0-147
	org.apache.zookeeper	zookeeper-contrib-loggraph	3.5.5.7.2.15.0-147
	org.apache.zookeeper	zookeeper-contrib-rest	3.5.5.7.2.15.0-147
	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.5.5.7.2.15.0-147
	org.apache.zookeeper	zookeeper-docs	3.5.5.7.2.15.0-147
	org.apache.zookeeper	zookeeper-jute	3.5.5.7.2.15.0-147
	org.apache.zookeeper	zookeeper-recipes-election	3.5.5.7.2.15.0-147
	org.apache.zookeeper	zookeeper-recipes-lock	3.5.5.7.2.15.0-147
	org.apache.zookeeper	zookeeper-recipes-queue	3.5.5.7.2.15.0-147

What's New In Cloudera Runtime 7.2.15

You must be aware of the additional functionalities and improvements to features of components in Cloudera Runtime 7.2.15. Learn how the new features and improvements benefit you.

What's New in Cruise Control

Learn about the new features of Cruise Control in Cloudera Runtime 7.2.15.

Configuration property for HTTP Strict Transport Security

There is a new configuration property for Cruise Control that enables Strict Transport Security header in the web server responses when SSL is enabled. By default, the configuration is enabled, and when TLS is enabled, Cruise Control sets the Strict Transport Security policy in the web server responses.

What's New in Hue

Learn about the new features of Hue in Cloudera Runtime 7.2.15.

Hue: The next generation SQL assistant

Hue packs the combined abilities of Data Analytics Studio (DAS) such as query optimization, query debugging framework, and rich query editor experience of Hue, making Hue the next generation SQL assistant on CDP. You can search Hive query history, view query details, visual explain plan, and DAG information, compare two queries, and download debug bundles for troubleshooting from the Job Browser page.

Ability to view Hive query history in Hue

A new service called Query Processor is added to the CDP stack as a dependency for Hue. It is used for indexing and retrieving Hive query history and query details. You can choose to turn on or turn off the Queries tab on the Hue Job Browser page by selecting or deselecting the Query Processor option in the Hue configurations in Cloudera Manager. For more information, see About Hue Query Processor.

What's New in Apache Kafka

Learn about the new features of Apache Kafka in Cloudera Runtime 7.2.15.

Enable JMX Authentication by default

JMX Authentication is now enabled by default for the Kafka service. Randomly generated passwords are now set for both the JMX monitor (read only access) and control (read and write access) users. The default passwords can be changed at any time using the Password of User with read-only Access to the JMX agent and the Password of user with read-write access to the JMX agent Kafka service properties. Additionally, JMX authentication can be turned off using the Enable Authenticated Communication with the JMX Agent property.

OAuth2 authentication available for Kafka

Oauth2 authentication support is added for the Kafka service. You can now configure Kafka brokers to authenticate clients using Oauth2. For more information, see OAuth2 authentication.

HSTS header is included by default in Kafka Connect REST API responses

Kafka Connect REST API responses now include the HSTS header by default.

Kafka load balancer support

The Kafka service can now be provided with a host of a load balancer that is used to balance connection bootstraps between multiple brokers. The host can be configured using the Kafka Broker Load Balancer Host property. Additionally, if a host is configured, the Kafka service configures a listener for accepting requests from the load balancer. This port is customizable using the Kafka Broker Load Balancer Listener Port property. Using these properties configures your Kafka service in a way that clients can connect to the brokers without encountering ticket mismatch issues in Kerberized environments or TLS/SSL hostname verification failures.

Importing Kafka entities into Atlas

Kafka topics and clients can now be imported into Atlas as entities (metadata) using a new action available for the Kafka service in Cloudera Manager. The new action is available at Kafka service>Actions>Import Kafka Topics Into Atlas. The action serves as a replacement/alternative for the kafka-import.sh tool. For more information, see Importing Kafka entities into Atlas.

Debezium Connector support

The following change data capture (CDC) connectors are added to Kafka Connect:

- Debezium MySQL Source
- Debezium Postgres Source
- Debezium SQL Server Source
- Debezium Oracle Source

Each of the connectors require CDP specific steps before they can be deployed. For more information, see Connectors.

Secure Kafka Connect

Kafka Connect is now generally available and can be used in production environments. This is the result of multiple changes, improvements, and new features related to Kafka Connect security including the following:

SPNEGO authentication for the Kafka Connect REST API

You can secure the Kafka Connect REST API by enabling SPNEGO authentication. If SPNEGO authentication is enabled, only users authenticated with Kerberos are able to access and use the REST API. Additionally, if Ranger authorization is enabled for the Kafka service, authenticated users will only be able perform the operations that they are authorized for. For more information, see Configuring SPNEGO Authentication and trusted proxies for the Kafka Connect REST API.

Kafka Connect Authorization model

An authorization model is introduced for Kafka Connect. Implementations are pluggable and it is up to the implementation how the capabilities of the model are utilized. The authorization model is implemented by default in Ranger. For more information about the model, see Kafka Connect authorization model. For more information about the Ranger integration of the model, see Kafka Connect Ranger integration.

Kafka Connect connector configurations can now be secured

A new feature called Kafka Connect Secrets Storage is introduced. This feature enables you to mark properties within connector configurations as a secret. If a property is marked as a secret, the feature stores and handles the value of that property in a secure manner. For more information, see Kafka Connect Secrets Storage.

Kafka Connect Connectors can be configured to override the JAAS, and restrict the usage of the Worker principal

Kafka Connect now allows users to force Connectors to override the JAAS configuration of the Kafka connection, and also forbid using the same Kerberos credentials as the Connect worker is using. For more information, see Configuring connector JAAS configuration and Kerberos principal overrides

Nexus allow list for Stateless NiFi Source and Sink connectors

A new configuration property, List Of Allowed Nexus Repository Urls, is introduced for the Kafka service. This property enables you to specify a list of allowed Nexus repositories that Kafka Connect connectors are allowed to connect to when fetching NiFi extensions. Configuring an allow list using the property can harden the security Kafka Connect deployment. For more information, see Configuring a Nexus repository allow list.

What's New in Apache Kudu

Learn about the new features of Kudu in Cloudera Runtime 7.2.15.

New tool to remove the dead tablet server

A new tool kudu tserver unregister is added to remove a dead tablet server from the cluster without restarting the masters. For more information, see KUDU-2915.

New column adding tool

A new tool kudu table add_column is added to add columns to existing tables. For more information, see KUDU-3339.

Tracking startup progress

It’s now possible to track startup progress of a Kudu server on the /startup page on the web UI. There are also metrics added to track the overall server startup progress as well as the processing of the log block containers and starting of the tablets. For more information, see KUDU-1959.

Improvements

- KUDU-3240: Client-side connection negotiation timeout is now configurable in the Java client
- KUDU-3328: The rebalancer tool now does not move replicas to tablet servers in maintenance mode
- KUDU-3340: It is now possible to disable compaction on a particular table.
- KUDU-3342: the kudu remote_replica list CLI tool now displays the data state and last status for a tablet replica.
- KUDU-3344: Kudu master cleans up metadata for deleted tables
- Table entity is now accessible in KuduWriteOperation in the C++ client, making understanding errors on the client side easier. For details, see KUDU-2623.
- Added pagination and search to the Tables page generated by the Kudu embedded Web server.
- The LZ4 library (Kudu uses it to compress various data on disk) has been upgraded to the 1.9.3 version to benefit from improved performance. For more information, see https://github.com/lz4/lz4/releases/tag/v1.9.3
- Fsync is now called on each modification of metadata files Kudu data directories are backed by XFS. The newly introduced –cmeta_fsync_override_on_xfs can be used to control this behavior.
- The log4j package used by the ranger-client plugin is upgraded to the 2.17.1 version.
- Run intra-location rebalancing in parallel: Intra-location rebalancing now runs concurrently at different locations for location-aware Kudu clusters. As a location-aware Kudu cluster automatically consists of non-intersecting groups of tablet servers, replicas within each location can be moved independently. Running intra-location rebalancing concurrently at every location can shorten the runtime of the rebalancer tool up to N times compared with running sequentially, where N is the number of locations defined in a Kudu cluster.

What's New in Schema Registry

Learn about the new features of Schema Registry in Cloudera Runtime 7.2.15.

Added OAuth support for Schema Registry client authentication

You can use OAuth2 JSON Web Token (JWT) in Schema Registry for authentication. Authorization continues to be implemented in Ranger, however, you can obtain the principal from a JWT token.

Added a findAllSchemas() method to the Schema Registry Client code

Provides a findAllSchemas() method which enumerates all schemas contained in the schema registry, returned as a list of SchemaMetadataInfo. This is useful if you only need to enumerate all schemas by name, without incurring the additional overhead of the findAggregatedSchemas() method.

Support for reading keys from JWK

Keys can be stored in JWK. The validation is done by matching with the kid property in JWT. If kid is not given then we match on the algorithm.

Added JWT validation filter

Added Servlet filter which checks if the incoming requests contain a valid authentication token.

SchemaRegistryClient gets token from OAuth Server with clientId/secret

Schema Registry Client can be configured to use OAuth2 authentication. The following parameters need to be added when creating a Schema Registry Client:

- "schema.registry.auth.type" = "oauth2" (default value is kerberos)
- "schema.registry.oauth.client.id" (ClientId for OAuth2 server)

- "schema.registry.oauth.secret" (Secret for OAuth2 server)
- "schema.registry.oauth.server.url" (REST API endpoint of OAuth2 server)

Support for RSA and HMAC certificates

Added support for JWT signed by either RSA or HMAC.

What's New in Cloudera Search

Learn about the new features of Cloudera Search in Cloudera Runtime 7.2.15.

Note: From Cloudera Runtime 7.2.15 and higher, new Cloudera Search features are listed under Apache Solr.

What's New in Apache Solr

Learn about the new features of Apache Solr in Cloudera Runtime 7.2.15.

- jQuery UI has been upgraded to version 1.13.0 due to CVEs.
- Apache Tika has been upgraded to version 2.3.0 due to CVEs.
- JDOM has been upgraded to version 2.0.6.1 due to CVEs.
- AWS SDK for Java has been upgraded to version 1.12.122 due to CVEs.

What's New in Sqoop

Learn what's new in the Apache Sqoop client in Cloudera Runtime 7.2.15.

To access the latest Sqoop documentation on Cloudera's documention 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 7.2.15.

Improvement in the Connect tab of the SMM UI

You can now deploy Kafka Connect connector configurations containing secret properties which will be stored in an encrypted storage (by default in Kafka). The deployed configuration will only contain references to these secrets. With this comes the need to mark properties as secret on the Streams Messaging Manager user interface so a new connector creation form is introduced, which supports it. You can import configurations and populate the form automatically.

Kafka Connect improvement

- In NiFi connectors you can now provide file path or URL for the flow.snapshot or alternatively you can upload it from file.

- You can now import Connector Configurations as a whole instead of adding individual configurations.
- Connector configuration validation errors are now correlated with individual config key.
- Sensitive properties are now hidden from the SMM UI and support is added to set properties as sensitive.

Partition dimension removal in SMM

The partition dimensions of the producer ("/api/v2/admin/metrics/aggregated/producers") and consumer ("/api/v2/admin/metrics/aggregated/groups") metrics are removed from the SMM cache, and are not exposed anymore through the API. This made the SMM memory footprint smaller, relieved some of the load from the metric store, and the network traffic became smaller. With this change, you get a cleaner, and easily readable API, and the UI is snappier, and faster than before.

The version of the /api/v1/admin/metrics/aggregated/* and /api/v1/admin/lineage/* endpoints have been changed to /api/v2/admin/aggregated and /api/v2/admin/lineage. With this change, the response objects are changed as well.

For the /lineage endpoints a common lineage response object is introduced in v2 as opposed to the specific (and different) objects in the experimental v1 endpoint.

For the /aggregated/* endpoints, the partition level metrics (that were in the wrappedPartitionMetrics field) are removed. Partition level metrics have been removed from the /aggregate/producers and /aggregated/producers/{producerClientId} but they are still available in the corresponding /metrics/producers and /metrics/producers/{producerId} endpoints.

Stateless Sink and Source should populate Key/Value Converters

SMM UI Connector Creation page now contains a default key/value converter to the StatelessNiFiSource or StatelessNiFiSink connectors.

Added API to enrich a sample configuration

Streams Messaging Manager API /connector-templates/config/enhance is added, which accepts a sample connector configuration and enhances it with the properties that are probably needed for that connector.

Add "emit.consumer.metrics" config to SMM CSD, and remove (now) unused SMON host/port configs

Removed "cm.metrics.service.monitor.host" and "cm.metrics.service.monitor.port" configurations from Streams Messaging Manager.

These no longer have to be configured as SMM automatically detects ServiceMonitor's location and emits the ConsumerGroup metrics into it.

Added "emit.consumer.metrics" configuration to Streams Messaging Manager.

In case this flag is disabled, Streams Messaging Manager does not emit historic ConsumerGroup metrics into ServiceMonitor, meaning historic metrics (for group Lag and CommittedOffset) would not be available for Groups in SMM. These metrics are used to populate the charts at the bottom of the ConsumerGroupDetail page, or accessed through the "api/v2/admin/metrics/consumers/group/{groupId}" REST API endpoint.

Increase SMM version to 2.3

SMM version is increased.

SMM UI should show the replication status tooltip

Streams Messaging Manager now shows tooltip for the replication status.

On the Overview page adjust the lineage information shown

On the Overview page, when a Producer or a Consumer is selected, an arrow points to the topic(s) it produced to or consumed from, instead of the partitions.

What's New in Streams Replication Manager

Learn about the new features of Streams Replication Manager in Cloudera Runtime 7.2.15.

New property Metric Reporting Period

A new property, Metric Reporting Period (metrics.period), is introduced for the SRM service. This property configures the frequency (in seconds) at which metrics should be reported.

SRM now creates all internal topics with correct configurations at startup

The internal topics used by SRM are now automatically created with correct configurations at startup. These are the metrics topics, the topics used by the srm-control tool, and the topics used by the SRM Service for service discovery. Additionally, SRM also verifies that the topics are created with correct configurations. If the topics are not configured as expected, SRM fails to start. This improvement fixes CDPD-31745.

Increase the default replication factor of internal topics to 3

The internal topics used by SRM are now created with a replication factor of 3 by default. As a result, SRM is now more resistant to host failures. Additionally, Cruise Control can now automatically heal SRM’s internal topics in the event of a single host failure.

SRM now waits for latest offset syncs and does not set the consumer offset into the future

The MirrorCheckpointConnector now checks the latest message in the offset sync topic at startup, and does not emit a checkpoint message until it has read from the beginning all the messages prior and including that last message.

As a part of this improvement, a new configuration property, emit.checkpoints.end.offset.protection is introduced. When this property is enabled, the MirrorCheckpointTask checks the end offset of the replicated topic prior to emitting a checkpoint, and limits the replicated offset to be maximum that value. With this behavior enabled, SRM no longer encounters an issue where in certain situations the replicated offset could be higher than the end offset of the replicated topic, producing a negative lag. The property is enabled by default, but can be configured using the Streams Replication Manager's Replication Configs property.

What's New in Apache Hadoop YARN

Learn about the new features of Hadoop YARN in Cloudera Runtime 7.2.15.

Queue priority

Setting queue priorities is now supported by the YARN Queue Manager UI. By setting queue priorities you can ensure that applications can access cluster resources. This is especially important in the case of Hive LLAP, long-running applications, and applications that require large containers.

For more information about the Dynamic Queue Scheduling feature, see Setting queue priorities.

Auto queue deletion at the queue level

Automatic removal of dynamically created child queues can be enabled and disabled not just globally, but also at the queue level. In addition, you can configure the queue expiration time.

For more information, see Deleting dynamically created child queues.

New YARN Queue Manager Overivew Page

In absolute mode the cluster resource capacity can be modified for the root queue using the YARN Queue Manager.

For more information, see Configuring the resource capacity of root queue in absolute mode.

Unaffected Components in this release

There are no new features for the following components in Cloudera Runtime 7.2.15.

- Apache Atlas
- Data Analytics Studio
- Apache HBase
- Apache Hadoop HDFS
- Apache Hive
- Apache Impala
- Apache Knox
- Apache Oozie
- Apache Phoenix
- Apache Ranger
- Apache Spark
- Apache ZooKeeper

Fixed Issues In Cloudera Runtime 7.2.15

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.

Fixed Issues in Atlas

Review the list of Apache Atlas issues that are resolved in Cloudera Runtime 7.2.15.

CDPD-39300: Upgraded jQuery-ui from version 1.12.1 to 1.13.1.

This issue is now resolved.

CDPD-35609: Jersey version upgrade to 1.19.4.

This issue is now resolved.

CDPD-34942: Currently in Atlas, code supports the above in "o3fs" ozone scheme. This patch provides improved support in "ofs" scheme, same as that of "o3fs".

This issue is now resolved.

CDPD-34903: Log4j-core dependency is removed from Atlas to avoid CVE.

This issue is now resolved.

CDPD-33113: Had dependency version problem, So upgraded the Babel version to latest.

This issue is now resolved.

CDPD-31937: Upgraded Underscore.js version to 1.13.1, due to CVE-2021-23358.

This issue is now resolved.

Apache patch information

- ATLAS-4580
- ATLAS-4569
- ATLAS-4570
- ATLAS-4556
- ATLAS-4546
- ATLAS-4549
- ATLAS-4529

- ATLAS-4527
- ATLAS-4499
- ATLAS-4492
- ATLAS-4522

Fixed Issues in Avro

There are no fixed issues for Avro in Cloudera Runtime 7.2.15.

Apache patch information

None

Fixed Issues in Cloud Connectors

Review the list of Cloud Connectors issues that are resolved in Cloudera Runtime 7.2.15.

CDPD-35142: Add Interface EtagSource to allow FileStatus subclasses to provide etags.

This issue is now resolved.

CDPD-35030: Fix multi object delete in S3A when number of objects is more than 1000.

This issue is now resolved.

CDPD-29477: ABFS: Partially obfuscate SAS object IDs in Logs.

This issue is now resolved.

Apache Patch Information

- HADOOP-18150
- HADOOP-18114
- HADOOP-18084
- HADOOP-16223
- HADOOP-18026
- HADOOP-17765
- HADOOP-18075
- HADOOP-17934
- HADOOP-17945
- HADOOP-14334
- HADOOP-17979
- HADOOP-17402
- HADOOP-16742
- HADOOP-17112
- HADOOP-18068
- HADOOP-17618

Fixed issues in Cruise Control

Review the list of Cruise Control issues that are resolved in Cloudera Runtime 7.2.15.

Support added for keystore and truststore types other than JKS

You are able to configure the keystore and truststore in Kafka brokers for Cruise Control Metrics Reporter. Previously, only the JKS type was supported for the SSL keystore and truststore.

Migrating Cruise Control to Log4j2

You are able to configure the keystore and truststore in Kafka brokers for Cruise Control Metrics Reporter. Previously, only the JKS type was supported for the SSL keystore and truststore.

Cruise Control fails to start after upgrade with Rack Aware Goal configured

You are able to configure the keystore and truststore in Kafka brokers for Cruise Control Metrics Reporter. Previously, only the JKS type was supported for the SSL keystore and truststore.

Fixed issues in Data Analytics Studio

There are no fixed issues for Data Analytics Studio in Cloudera Runtime 7.2.15.

Fixed Issues in Apache Hadoop

Review the list of Hadoop issues that are resolved in Cloudera Runtime 7.2.15.

CDPD-34911: Harden FileUtil.unTar(File inFile, File untarDir).

This issue is now resolved.

CDPD-35233: LazyPersist Overwrite fails in direct write mode.

This issue is now resolved.

CDPD-35138: Add unresolved endpoint value to UnknownHostException.

This issue is now resolved.

Apache Patch Information

- HADOOP-18136
- HADOOP-17328
- HADOOP-13500
- HADOOP-17152
- HADOOP-17929
- HADOOP-17126
- HADOOP-17930
- HADOOP-17947
- HADOOP-17897
- HADOOP-17837
- HADOOP-16827
- HADOOP-17587
- HADOOP-17610

Fixed Issues in HBase

Review the list of HBase issues that are resolved in Cloudera Runtime 7.2.15.

Apache Patch Information

- HBASE-26880
- HBASE-26469
- HBASE-26512
- HBASE-26896
- HBASE-26838
- HBASE-26885
- HBASE-26832

- HBASE-26872
- HBASE-26675
- HBASE-25709
- HBASE-23303
- HBASE-26576
- HBASE-26712
- HBASE-26742
- HBASE-26434
- HBASE-26688
- HBASE-26741
- HBASE-26747
- HBASE-26713
- HBASE-26679
- HBASE-26662
- HBASE-26657
- HBASE-26643
- HBASE-26646
- HBASE-26590
- HBASE-26629
- HBASE-26625
- HBASE-26615
- HBASE-26609
- HBASE-26598
- HBASE-26579
- HBASE-26488
- HBASE-26340
- HBASE-26271
- HBASE-26530
- HBASE-24163
- HBASE-26107
- HBASE-26168
- HBASE-26124
- HBASE-26525
- HBASE-26527
- HBASE-26462
- HBASE-26533
- HBASE-26517
- HBASE-26468
- HBASE-25905
- HBASE-26485
- HBASE-26476
- HBASE-26450
- HBASE-26465
- HBASE-26482
- HBASE-26475
- HBASE-26470
- HBASE-26429
- HBASE-26311
- HBASE-26309
- HBASE-26406

- HBASE-26384
- HBASE-26394
- HBASE-26327
- HBASE-26390
- HBASE-26190
- HBASE-26308
- HBASE-26382
- HBASE-26385
- HBASE-26190
- HBASE-26371
- HBASE-26312
- HBASE-26339
- HBASE-26238
- HBASE-26297
- HBASE-26274
- HBASE-26261
- HBASE-26281
- HBASE-26273
- HBASE-26276
- HBASE-26255
- HBASE-26229

Fixed Issues in HDFS

Review the list of HDFS issues that are resolved in Cloudera Runtime 7.2.15.

Apache Patch Information

- HDFS-14655

Fixed Issues in Apache Hive

Review the list of Hive issues that are resolved in Cloudera Runtime 7.2.15.

Apache Patch Information

- HIVE-26029
- HIVE-25686
- HIVE-21498

Fixed Issues in Hive Warehouse Connector

There are no fixed issues for HWC in Cloudera Runtime 7.2.15.

Fixed Issues in Hue

There are no fixed issues for Hue in Cloudera Runtime 7.2.15.

Fixed Issues in Apache Impala

Review the list of Impala issues that are resolved in Cloudera Runtime 7.2.15.

CDPD-35202: Impala logs too much error message when the log file symlink is not found.

This issue is now resolved.

Technical Service Bulletin

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

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

Fixed Issues in Apache Kafka

Review the list of Apache Kafka issues that are resolved in Cloudera Runtime 7.2.15.

CDPD-29058: Migrate to log4j2 due to log4j1 end of life

Kafka is migrated and uses log4j2 as a logging library. Additionally, log4j1 dependencies are removed with the exception of the Log4jAppender. Although the appender remains available, Cloudera recommends that you use the log4j2 implementation of the appender that is available in the log4j2 project.

CDPD-29307: Kafka keystore and truststore type is not configured for Cruise Control metrics reporter

The keystore and truststore types are now correctly supported by the Cruise Control metrics reporter in the Kafka broker.

OPSAPS-62548: TopicMetrics get deleted from Cloudera Manager during restart or Kafka partition reassignment

KafkaTopicMetrics are no longer deleted from the ServiceMonitor's Time-series database during a Kafka restart or a partition leader change.

Apache Patch Information

- KAFKA-13443: Kafka broker exits when OAuth enabled and certain configuration not specified
- KAFKA-13445: Add ECDSA test for JWT validation
- KAFKA-13444: Fix OAuthCompatibilityTool help and add SSL options
- KAFKA-13446: Remove JWT access token from logs
- KAFKA-13202: KIP-768: Extend SASL/OAUTHBEARER with Support for OIDC

Fixed Issues in Apache Knox

Review the list of Knox issues that are resolved in Cloudera Runtime 7.2.15.

Apache patch information

Apache patches in this release.

- KNOX-2735
- KNOX-2731
- KNOX-2710
- KNOX-2714
- KNOX-2713
- KNOX-2712
- KNOX-2531

- KNOX-2708
- KNOX-2149
- KNOX-2705
- KNOX-2704
- KNOX-2698
- KNOX-2693
- KNOX-2690
- KNOX-2699

Fixed Issues in Apache Kudu

There are no fixed issues for Kudu in Cloudera Runtime 7.2.15.

Unpartitioned table cannot be copied

Copying unpartitioned tables using the kudu table copy CLI tool is now supported.

Failing rebalancer

Fixed a bug where the rebalancer failed when trying to decommission a tablet server in a location-aware Kudu cluster. For more information, see KUDU-3346.

Java client: malformed tabler server entry when rehydrated from a scan token

Fixed a bug in the Java client where a malformed tablet server ID in the scan token causes connection failures and timeouts in some cases. For more information, see KUDU-3349.

Unpartitioned table cannot be copied

Fixed a bug which could lead to exhaustion of the address space for the outgoing connections on a busy Kudu cluster. For more information, see KUDU-3352.

Fixed Issues in Livy

Review the list of Livy issues that are resolved in Cloudera Runtime 7.2.15.

CDPD-19276: The HTTP headers Strict-Transport-Security (HSTS) and Content-Security-Policy (CSP) are populated in Livy and Livy for Spark 3 by default and their values can be customized via a configuration setting.

This issue is now resolved.

Apache patch information

None

Fixed Issues in Apache Oozie

Review the list of Oozie issues that are resolved in Cloudera Runtime 7.2.15.

CDPD-34649: The following issue was fixed: fs:fileSize returned zero when the path was not normalized even if the file size was not zero.

This issue is resolved.

CDPD-34475: Fixed a possible dead-lock in SignalXCommand.

This issue is resolved.

CDPD-33993: Oozie will not turn to HDFS NameNode to "resolve" the share lib's file entries one by one before starts the Yarn application. If somehow this causes fault then this functionality can be turned off via a Cloudera Manager Safety Valve: oozie.classpathutils.resolve=true.

This issue is resolved.

CDPD-16552: Users are now able to configure the Action directory path. For more information please see the Action Configuration section in the Oozie documentation.

This issue is resolved.

CDPD-14956: When having a log4j.properties file under sharelib, Oozie set a invalid environment variable which failed the container.

This issue is resolved.

Apache patch information

- OOZIE-3658
- OOZIE-3524
- OOZIE-3646
- OOZIE-3545
- OOZIE-3606
- OOZIE-3602

Fixed Issues in Ozone

There are no fixed issues for Ozone in Cloudera Runtime 7.2.15.

Apache patch information

None

Fixed Issues in Phoenix

Review the list of Phoenix issues that are resolved in Cloudera Runtime 7.2.15.

CDPD-35717: CALCITE-903 has introduced a transaprent reconnection feature, which will open a new server-side connection in case it is expired from the server side connection cache. While this is convinient for most read-only analytical workloads, this can cause a number a problems, including data loss for transactional connections. This patch disables the transparent reconnect feature by default, and adds the transparent_reconnection property, which re-enables it when set to true.

This issue is resolved.

CDPD-35652: PreparedStatement#getMetaData() no longer fails on parametrized "select next ? values" sequence operations. This also fixes a problem where parametrized sequence operations didn't work via Phoenix Query Server.

This issue is resolved.

CDPD-35391: The phoenix-sqlline and phoenix-sqlline-thin console SQL clients now work correctly on PowerPC architecture.

This issue is resolved.

CDPD-35250: Phoenix versions older than 4.14 set the KEEP_DELETED_CELLS and VERSIONS properties on the SYSTEM.STATS and SYSTEM.LOG table, which caused performance problems. This problem has been fixed for fresh installs in Phoenix 4.14 and 5.0, but on Phoenix installations being upgraded from pre 4.14 versions, the properties were not removed. Now those properties are removed when performing an upgrade from and older version.

This issue is resolved.

Apache Patch Information

- CALCITE-5009
- PHOENIX-6665

- PHOENIX-6661
- PHOENIX-6645
- PHOENIX-6646
- PHOENIX-5894
- PHOENIX-6579
- PHOENIX-6611
- PHOENIX-6632
- PHOENIX-6654

Fixed Issues in Parquet

There are no fixed issues for Parquet in Cloudera Runtime 7.2.15.

Apache Patch Information

None

Fixed Issues in Apache Ranger

Review the list of Ranger issues that are resolved in Cloudera Runtime 7.2.15.

CDPD-39317: Updated atlas default audit filter to avoid auditing for atlas read-entity by nifi service user.

This issue is now resolved.

CDPD-38668: S3 plugin reports the result of the evaluation of S3 access policies differently than HDFS plugin. The fix is to change the handling of the result of base-plugin's policy evaluation result so that it works in all cases.

This issue is now resolved.

CDPD-36327: Enabled HDFS or cloud storage auditing for the Kafka Connect Ranger plugin.

This issue is now resolved.

CDPD-35742: Change display messages in policy form items.

This issue is now resolved.

CDPD-35631: Fixed role update operation issue for role admin user (A non admin user should be able to update the role if user is role admin).

This issue is now resolved.

CDPD-35204: Flagged based enhancement enables to transform/update username using user-mapping file provided at the time sentry migration using authzmigrator tool.

This issue is now resolved.

CDPD-35073: Upgrade jquery-ui 1.12 to 1.13.0+ due to CVEs.

This issue is now resolved.

CDPD-35073: Flagged based enhancement enables to create S3 policy for Hive warehouse location at the time of sentry migration.

This issue is now resolved.

CDPD-34762: User/group/tags/resource attributes should be easily accessible in condition expressions, with expressions like: USER.state == 'CA' UG['test'].dept == 'MKTG' REQ.accessType == 'SELECT' RES.database == 'hr' RES.table == 'employee' TAG._type == 'PII' TAG.attr1 == 'value1' TAGS.PII.attr1 == 'value1' TNAMES.length == 2 TNAMES.indexOf('PCI') != -1

This issue is now resolved.

CDPD-34750: This change is to add the support for retry for policies download, ugsync, tagsync.

This issue is now resolved.

CDPD-34723: Policy engine evaluates policies in the following order: priority, has-deny, has-no-deny. When multiple policies have same priority/has-deny/has-no-deny, the ordering is not deterministic. This doesn't impact the result for access policies - as all denies will be evaluated before allows. However, the result for masking/row-filter can vary when multiple policies exists for a given resource, and these policies define different mask/filter for a given user/group/role.

Given name of a policy is unique within a service, using policy name as the secondary sorting key will result in deterministic evaluation order. This issue is now resolved.

CDPD-34057: Updated Ranger db setup to support Mysql DB versions from and above version 8.0.

This issue is now resolved.

CDPD-34023: Chmod and Chown will honor the ranger policy in both with fallback enabled as well as disabled. Workaround is to have the parent directory RX permission in HDFS for the failing folders/files.

This issue is now resolved.

CDPD-33058: Resolve UI side regression for rendering resources.

This issue is now resolved.

CDPD-32975: Storm library version in Ranger upgraded to fix the CVE.

This issue is now resolved.

CDPD-32974: kylin library version in Ranger upgraded to fix the CVE.

This issue is now resolved.

CDPD-32879: Added a config "ranger-rms.max.requested.notifications" to limit the size of requested notifications during the delta-sync. Setting the config value < 1 or > 50000; the default value will be treated as maxRequestedNotifications=50000. The default value for MAX_REQUESTED_NOTIFICATIONS is 50000. This fix includes the bugs: handleDeltaSync loop runs infinite when it tries to fetch notifications in batch and Full-sync does not reset last_known_version=-1 in x_rms_mapping_provider table.

This issue is now resolved.

CDPD-32874: Improvement in load permission edit page with more number of users and groups data.Added lazy loading for that.

This issue is now resolved.

CDPD-31780: This change is to integrate/certify the Ranger DB KMS with GCP.

This issue is now resolved.

CDPD-31358: During the upgrade this change will update the existing solr policies to fit with new resource and permission types.

This issue is now resolved.

CDPD-31357: Click on the policy resource field that time all available resource options are listed down.

This issue is now resolved.

CDPD-31127: Upgrade netty to 4.1.68+.

This issue is now resolved.

CDPD-28944: Newly created tag/value in policy resources field displayed with "Create" tag.

This issue is now resolved.

CDPD-28752: Provided sorting on specific columns on policy and on audit listing page.

This issue is now resolved.

CDPD-26846: User who wishes to configure custom cipher suite can add a custom property ranger.usersync.https.ssl.enabled.cipher.suites with comma separated value of required ciphers and restart Ranger Usersync.

This issue is now resolved.

CDPD-26334: Fixed tooltip hit for a better user understanding.

This issue is now resolved.

CDPD-22073: While adding value to policy resources, if resource lookup fails then it doesn't show any indication to the user for resource lookup fail.

When this issue occurs, a notification is displayed. This issue is now resolved.

CDPD-16420: Added pause/play button to notification popup.

This issue is now resolved.

Apache Patch Information

- RANGER-3725
- RANGER-3691
- RANGER-3699
- RANGER-3690
- RANGER-3675
- RANGER-3665
- RANGER-3659
- RANGER-3509
- RANGER-3290
- RANGER-3661
- RANGER-3600
- RANGER-3642
- RANGER-3617
- RANGER-3609
- RANGER-3606
- RANGER-3605
- RANGER-3586
- RANGER-3567
- RANGER-3610
- RANGER-3550
- RANGER-3508
- RANGER-3565
- RANGER-3620
- RANGER-3526
- RANGER-3613
- RANGER-3647
- RANGER-3592
- RANGER-3591
- RANGER-3577
- RANGER-3638
- RANGER-3603
- RANGER-3568
- RANGER-3552
- RANGER-3539
- RANGER-3569
- RANGER-3629
- RANGER-3649
- RANGER-3520
- RANGER-3485
- RANGER-3459
- RANGER-3427

- RANGER-3403
- RANGER-3667
- RANGER-3533
- RANGER-3673
- RANGER-3660
- RANGER-3630
- RANGER-3594
- RANGER-3576
- RANGER-3666
- RANGER-3662

Fixed issues in RAZ

Review the list of RAZ issues that are resolved in Cloudera Runtime 7.2.15.

CDPD-35036: Fixing signature catching in RAZ S3 by increasing expiry time check interval.

This issue is resolved.

CDPD-31371: Added RAZ policy for RAZ Azure datalake auto backup/ restore.

This issue is resolved.

Fixed issues in Spark Atlas Connector

Review the list of Spark Atlas Connector issues that are resolved in Cloudera Runtime 7.2.15.

CDPD-34679: Input file names containing spaces caused URISyntaxException to be thrown.

This issue is resolved.

Fixed Issues in Schema Registry

Review the list of Schema Registry issues that are resolved in Cloudera Runtime 7.2.15.

CDPD-35983: Unique constraint violation on load balanced Schema Registry cluster startup

A concurrency issue in a multi-node Schema Registry setup is fixed where more nodes tried to initialize database state at the same time causing some of them to fail.

CDPD-35469: Schema Registry responds with Internal Server Error when adding more schemas than defined in offset range

Schema Registry responds with HTTP 409 response instead of HTTP 500 response while trying to add more schemas than defined in offset range.

CDPD-33908: Remove or Upgrade Spring framework to 5.3.14+/5.2.19 due to CVE-2021-22060

Removed Spring dependencies from Schema Registry because they were not used at all.

CDPD-32192: First start failed for Schema Registry, with oracle DB, migration failed at CREATE TABLE "atlas_events"

Fixed v009__create_registry_audit.sql to have create index refer to the lower case "atlas_events" object (the table).

Made the script rerunnable since the table was already created where the script had already run.

CDPD-31881: Schema Registry L1 test fails with socket timeout

When more than one instance of Schema Registry is running on the same DB, "concurrent update" exceptions might have appeared in the Schema Registry log regarding changes to be sent to Atlas.

Fixed Issues in Cloudera Search

Review the list of Cloudera Search issues that are resolved in Cloudera Runtime 7.2.15.

Note: From Cloudera Runtime 7.2.15 and higher, new Cloudera Search fixed issues are listed under Apache Solr.

Fixed Issues in Apache Solr

Review the list of Solr issues that are resolved in Cloudera Runtime 7.2.15.

Number of warning messages in Solr logs reduced (backport of SOLR-11623)

Every request handler now implement PermissionNameProvider interface and provide "ALL" permission type if no authorization checks are necessary.

Support for HSTS Security Protocol added (Backport of SOLR-15578)

Manually editing the jetty.xml file as a workaround to enable HSTS is no longer necessary.

CDPD-24003: The Solr admin UI is only accessible with full solr_admin permission in Ranger

Full solr_admin permission is no longer required in Ranger to access the Solr Admin UI.

Technical Service Bulletins

TSB-847: CVE-2025-30065 Apache Parquet vulnerability

On April 1, 2025, a critical vulnerability in the parquet-avro module of Apache Parquet (CVE-2025-30065, CVSS score 10.0) was announced.

Remediation for affected versions

The Cloudera Search release patched through the CDP updates for the public cloud and private cloud base.

Note: Cloudera will not provide remediation options for unsupported versions, and has not tested mitigations on unsupported versions. Customers are advised to upgrade to a supported product version. For more information, refer to the Support Lifecycle Policy page.

Vulnerability details

Exploiting this vulnerability is only possible by modifying the accepted schema used for translating Parquet files and subsequently submitting a specifically crafted malicious file.

Schema parsing in the parquet-avro module of Apache Parquet 1.15.0 and previous versions allows bad actors to execute arbitrary code. Attackers may be able to modify unexpected objects or data that was assumed to be safe from modification. Deserialized data or code could be modified without using the provided accessor functions, or unexpected functions could be invoked.

Deserialization vulnerabilities most commonly lead to undefined behavior, such as memory modification or remote code execution.

Action required - Mitigation for affected Cloudera products:

Until the upgrade with Apache Parquet 1.15.1 or higher is available:

1. Utilize a File Integrity Monitoring (FIM) solution. This allows administrators to monitor files at the filesystem level and receive alerts on any unexpected or suspicious activity in the schema configuration.
2. Monitor network activity for any transmission of Parquet files, and alert on any unexpected activity.
3. Be cautious with Parquet files from unknown or untrusted sources. If possible, do not process files with uncertain origin or that came from outside the organization.

4. Ensure that only authorized users have access to endpoints that ingest Parquet files.

For the latest update on this issue see the corresponding Knowledge Article: TSB 2025-847: Critical Apache Parquet vulnerability CVE-2025-30065

Fixed Issues in Spark

Review the list of Spark issues that are resolved in Cloudera Runtime 7.2.15.

CDPD-30076: Fix issue where TableSplit always returns size of "0" under 1MB, even if data is present.

This issue is resolved.

Apache patch information

MAPREDUCE-7341

Fixed Issues in Apache Sqoop

There are no fixed issues for Sqoop in Cloudera Runtime 7.2.15.

Apache patch information

None

Fixed Issues in Streams Messaging Manager

Review the list of Streams Messaging Manager issues that are resolved in Cloudera Runtime 7.2.15.

CDPD-33770: On the topics details page selecting a custom timestamp is broken

Fixed SMM REST throwing an internal server error when custom timestamps are provided while calling the "/api/v2{or v1}/admin/replication-stats" endpoint, or when a custom time period is provided on the ProducerDetail page in SMM UI.

CDPD-33011: Selecting a consumer with no producers should show 0 producers in the filter panel

On the overview page in the filter panel, when a consumer is selected that has no producers associated, the number of producers will be shown to be 0 of T, where T is the total number of producers.

CDPD-32936: Selecting a producer with no consumers should show 0 consumers in the filter panel

On the overview page in the filter panel, when a producer is selected that has no consumers associated, the number of consumers will be shown to be 0 of T, where T is the total number of consumers.

CDPD-29403: When editing the alert, the topic can be chosen for the replication status

Fixed the topic selection dropdown status in the alert editor after various UI events.

OPSAPS-63017: The Kafka Connect tab is missing from the SMM UI

The Kafka Connect tab is now correctly displayed if Kafka Connect is provisioned on the cluster.

OPSAPS-62548: TopicMetrics get deleted from CM during restart or kafka partition reassignment

KafkaTopicMetrics accidentally gets deleted from ServiceMonitor's Timeseries database during a Kafka restart or partition leader change.

Fixed Issues in Streams Replication Manager

Review the list of Streams Replication Manager issues that are resolved in Cloudera Runtime 7.2.15.

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

SRM now creates all internal topics explicitly. SRM also verifies the essential configurations of internal topics at startup, and fails if the topics do not meet the required configurations.

OPSAPS-63104: The automatically generated password for co-located services is invalid

SRM Service Basic Authentication would not work with the default, random generated password. SRM Service Basic Authentication default password is now identical on all SRM Service role instances.

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

SRM uses the correct ssl.keystore.key configuration when a Kafka External Account specifies the keystore.

Fixed Issues in Apache YARN

Review the list of Apache Hadoop YARN issues that are resolved in Cloudera Runtime 7.2.15

COMPX-5255: Can't enable "auto queue creation" for "root" in weight mode

Enabling the dynamic child queue creation feature for root queue is now possible.

COMPX-7292: Add option to disable auto creation of a queue in weight mode

Previously if you enabled the dynamic child creation feature for a queue you could not disabled it; you had to remove the parent queue and then recreate it. This fix enables you to disable the auto child creation feature for a queue.

COMPX-7594: Minimum User Limit field fractional percentages inconsistency

The Minimum User Limit field previously accepted fractional percentages. This issue has been fixed and the field does not accept francional percentages anymore.

COMPX-7619: Placement Rules "View" button should be available in read-only mode

This fix adds the View options for placement rules for non-admin users when read-only mode is allowed for them in YARN Queue Manager UI.

COMPX-7822: User Limit Factor property should accept -1

The User Limit Factor property now accepts -1 as a value.

COMPX-8360: Yarn Queue Manager UI wipes out all Placement Rules when an invalid rule is added

When an invalid placement rules logic was created, the previously created valid placement rules were removed by the YARN Queue Manager UI, while the configuration file did not remove them. With this fix, the YARN Queue Manager UI does not remove the previously created valid placement rules anymore.

Apache patch information

- YARN-6862
- YARN-10503
- YARN-10870

Fixed Issues in Zeppelin

There are no fixed issues for Zeppelin in Cloudera Runtime 7.2.15.

Apache patch information

None

Fixed Issues in Apache ZooKeeper

Review the list of ZooKeeper issues that are resolved in Cloudera Runtime 7.2.15.

Apache Patch Information

- ZOOKEEPER-3652: Synchronize ClientCnxn outgoing queue flush on a stable internal value
- ZOOKEEPER-4477: Single Kerberos ticket renewal failure can prevent all future renewals since Java 9

Fixed Issues In Cloudera Runtime 7.2.15.1

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.1.

CFM-2775: Added download flow definition with external services menu and fixed regressions

This issue is now resolved.

Fixed Issues In Cloudera Runtime 7.2.15.2

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.2.

The following issues are resolved:

- HOTREQ-968 Large file upload to WebHDFS causes OOM in Knox
- HOTREQ-1000 Mutually exclusive filter params in the HadoopGroupProvider identity-assertion provider
- HOTREQ-1001 NPE occurred while getting service discovery types
- HOTREQ-964 Release: HIVE-25574: Replace clob with varchar when storing creation metadata
- HOTREQ-1003 Optimisations on COD for ABFS support
- HOTREQ-1026 KnoxToken doAs broken with HadoopAuthFilter

Known Issue: CDPD-43048

SecureBulkLoad feature does not work on public cloud when using S3 with HBOSS and Raz.

Workaround: After 7.2.15.2, the next ODX version will be released. This enables SFT and disable HBOSS on CDPD 7.2.15.2. This will fix SecureBulkload.

Fixed Issues In Cloudera Runtime 7.2.15.3

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.3.

The following issues are resolved:

- HOTREQ-1052 Oozie Actions which rely on "action-data.seq" failed intermittently
- HOTREQ-1091 Casting invalid dates does not produce NULL

Known Issue: CDPD-43048

SecureBulkLoad feature does not work on public cloud when using S3 with HBOSS and Raz.

Workaround: After 7.2.15.2, the next ODX version will be released. This enables SFT and disable HBOSS on CDPD 7.2.15.2. This will fix SecureBulkload.

Fixed Issues In Cloudera Runtime 7.2.15.4

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.4.

The following issues are resolved:

- HOTREQ-1102 Hive compaction state turns "ready for cleaning
- HOTREQ-1051 Query based compaction fails in Public Cloud
- HOTREQ-1036 Bug Fix for SPARK-39083
- HOTREQ-1070 Backport for SPARK-38318
- HOTREQ-1114 Hue does not work with medium duty DL because IDBroker config has comma separated URLs

Known Issue: CDPD-43048

SecureBulkLoad feature does not work on public cloud when using S3 with HBOSS and Raz.

Workaround: After 7.2.15.2, the next ODX version will be released. This enables SFT and disable HBOSS on CDPD 7.2.15.2. This will fix SecureBulkload.

Technical Service Bulletins

TSB 2022-614: Kafka policy "user auto-creation" does not work in Ranger in CDP Public Cloud 7.2.15

For the latest update on this issue see the corresponding Knowledge article: TSB 2022-614: Kafka policy "user auto-creation" does not work in Ranger in CDP Public Cloud 7.2.15

Fixed Issues In Cloudera Runtime 7.2.15.5

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.5.

The following issues are resolved:

- HOTREQ-1171 - Hotfix request for fixing critical CVEs (CDPD-44986)
- HOTREQ-1178 - HUE Oozie workflow rerun fails.
- HOTREQ-1183 - Incorrect case evaluation for Parquet based table
- HOTREQ-1180 - Backport HIVE-25275 to public cloud 7.2.15
- HOTREQ-1202 - Implement support for preventing incompatible log4j classes to be loaded in Sqoop
- HOTREQ-1161 - CFM - NIFI nodes are getting disconnected frequently

Known Issue: CDPD-43048

SecureBulkLoad feature does not work on public cloud when using S3 with HBOSS and Raz.

Workaround: After 7.2.15.2, the next ODX version will be released. This enables SFT and disable HBOSS on CDPD 7.2.15.2. This will fix SecureBulkload.

Technical Service Bulletins

TSB 2022-614: Kafka policy "user auto-creation" does not work in Ranger in CDP Public Cloud 7.2.15

For the latest update on this issue see the corresponding Knowledge article: TSB 2022-614: Kafka policy "user auto-creation" does not work in Ranger in CDP Public Cloud 7.2.15

Fixed Issues In Cloudera Runtime 7.2.15.6

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.6.

The following issue is resolved:

- CDPD-46653 - File not found exception when StoreFileScanner is refreshed after flush followed by compaction.

Known Issue: CDPD-43048

SecureBulkLoad feature does not work on public cloud when using S3 with HBOSS and Raz.

Workaround: After 7.2.15.2, the next ODX version will be released. This enables SFT and disable HBOSS on CDPD 7.2.15.2. This will fix SecureBulkload.

Technical Service Bulletins

TSB 2022-614: Kafka policy "user auto-creation" does not work in Ranger in CDP Public Cloud 7.2.15

For the latest update on this issue see the corresponding Knowledge article: TSB 2022-614: Kafka policy "user auto-creation" does not work in Ranger in CDP Public Cloud 7.2.15

Fixed Issues In Cloudera Runtime 7.2.15.7

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.7.

The following issues are resolved:

- CDPD-46871 - Improve CleanerChore delete order with back reference links
- CDPD-46655 - File not found exception on StoreScanners upon a memstore flush

Known Issue: CDPD-43048

SecureBulkLoad feature does not work on public cloud when using S3 with HBOSS and Raz.

Workaround: After 7.2.15.2, the next ODX version will be released. This enables SFT and disable HBOSS on CDPD 7.2.15.2. This will fix SecureBulkload.

Technical Service Bulletins

TSB 2022-614: Kafka policy "user auto-creation" does not work in Ranger in CDP Public Cloud 7.2.15

For the latest update on this issue see the corresponding Knowledge article: TSB 2022-614: Kafka policy "user auto-creation" does not work in Ranger in CDP Public Cloud 7.2.15

Fixed Issues In Cloudera Runtime 7.2.15.8

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.8.

CDH

The following issues are resolved:

- HOTREQ-1245 - Ranger RAZ - Azure parent directory delete issue
- HOTREQ-1241 - Patch for SPARK-32380 (Spark 3) - sparksql cannot access hive table while data in HBase
- HOTREQ-1244 - HOTFIX Request for CDPD-46957 - Hue - Export feature not available from search bar in CDP
- HOTREQ-1222 - ABFS: Disable readAhead for 7.2.12 onwards PCR versions
- HOTREQ-1201 - HADOOP-18476 Abfs and S3A FileContext bindings to close wrapped filesystems in finalizer
- HOTREQ-1264 - HotFix for CDPD-39592 & CDPD-29225
  - CDPD-39592 - Due to the changes made in azure conf.py via PR-2396, not able to access Azure FB.
  - CDPD-29225 - Hue does not work with medium duty DL because IDBroker config has comma separated URLs

- HOTREQ-1275 - Hotfix for - IMPALA-11751 - Crash in processing partition columns of Avro table with MT_DOP>1
- HOTREQ-1274 - HOTFIX REQ for issue HUE with KNOX unable to open workflow links in new tab with right click

CFM

The following issues are resolved:

- HOTREQ-1223 - CFM - HTTP proxy password disappears every time cluster is restarted
- HOTREQ-1253 - Ranger-NiFi repo connectivity gets 403 error
- HOTREQ-1252 - Nifi nodes corruptions - Trying to start the Data hub cluster the underline NIFI service not starting properly.
- HOTREQ-1251 - Improve the memory management of stateless flow

Cloudera Manager

The following issues are resolved:

- HOTREQ-1258 - Request for include OPSAPS-64187 and OPSAPS-65242
  - OPSAPS-64187 - Cloudera Manager Event server does not clean up old events
  - OPSAPS-65242 - Alert Server event cleanup has further problems

CVE

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

Known Issue: CDPD-43048

SecureBulkLoad feature does not work on public cloud when using S3 with HBOSS and Raz.

Workaround: After 7.2.15.2, the next ODX version will be released. This enables SFT and disable HBOSS on CDPD 7.2.15.2. This will fix SecureBulkload.

Technical Service Bulletins

TSB 2022-614: Kafka policy "user auto-creation" does not work in Ranger in CDP Public Cloud 7.2.15

For the latest update on this issue see the corresponding Knowledge article: TSB 2022-614: Kafka policy "user auto-creation" does not work in Ranger in CDP Public Cloud 7.2.15

TSB 2023-644: Microsoft Azure parent directory deletion

For the latest update on this issue, see the corresponding Knowledge Base article: TSB 2023-644: Microsoft Azure parent directory deletion.

Fixed Issues In Cloudera Runtime 7.2.15.10

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.10.

The following issues are resolved:

- HOTREQ-1363 - Patch 7.2.15 with HIVE-25502 and release to customers (ENGESC-19275)
- TSB 2023-653: Cleaner causes data loss when processing an aborted dynamic partitioning transaction

Fixed Issues In Cloudera Runtime 7.2.15.11

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.11.

The following issues are resolved:

CDH

- HOTREQ-1330 add a way to reenable abfs readahead
- HOTREQ-1363 Patch 7.2.15 with HIVE-25502 and release to customers (ENGESC-19275)
- HOTREQ-1287 Wrong results for partitioned Parquet table when files contain partition column
- HOTREQ-1320 HOTFIX for Bug - Add delegation token support for long running spark job
- HOTREQ-1315 Upgrade node.js due to CVE-2022-35255, CVE-2022-43548 and CVE-2022-32212
- HOTREQ-1344 Hot fix JIRA CDPD-47077 for Public Cloud 7.2.15
- HOTREQ-1334 S3 Copy Optimization

CFM

- HOTREQ-1282 CFM - Parameter context inheritence fail during startup
- HOTREQ-1263 CFM - Nifi HWXSchemaRegistry Controller service connection on CDP Public cloud

Fixed Issues In Cloudera Runtime 7.2.15.12

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.12.

The following issues are resolved:

HOTREQ-1408: [Hue][ABFS] List root directories in RAZ enabled environment.

Fixed Issues In Cloudera Runtime 7.2.15.100

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.100.

CDH

- HOTREQ-1422 Need HIVE-26779 on Public Cloud runtime 7.2.15.10-1
- HOTREQ-1432 HIVE-23806 on 7.2.15
- HOTREQ-1433 Backport CDPD-55226 / HADOOP-18705 to 7.2.15
- HOTREQ-1445 Need SPARK-39376 On Public Cloud Runtime 7.2.15.10-100 (PCRR-268)

CFM

- HOTREQ-1423 ExecuteScript processor not supporting Module Directory for python
- HOTREQ-1374 Fix nifiregistry to use proper jdbc drivers

Fixed Issues In Cloudera Runtime 7.2.15.200

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.200.

CDH

- HOTREQ-1468 keytab does not exist error - ENGESC-20539
- HOTREQ-1461 HOTFIX for CDPD-48847
- HOTREQ-1443 HOTFIX for ENGESC-20387 (Yarn cluster overview pg refresh button issue)

CFM

- HOTREQ-1481 Ship NIFI-11744 security fix for 7.2.17 (CFM-2.2.7), 7.2.16 (CFM-2.2.6) and 7.2.15 (CFM-2.2.5)
- HOTREQ-1372 CaptureChangeMySQL processor fixes
- HOTREQ-1458 Ship NIFI-11614 security fix for 7.2.16 (CFM-2.2.6) and 7.2.15 (CFM-2.2.5)
- HOTREQ-1459 SHIP NIFI-11653 Security fix for 7.2.17 (CFM-2.2.7), 7.2.16 (CFM-2.2.6) and 7.2.15 (CFM-2.2.5)

Fixed Issues In Cloudera Runtime 7.2.15.300

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.300.

CDH

- HOTREQ-1515: Need HIVE-13288 on top of CDP 7.2.16
- HOTREQ-1429: Request to backport CDPD-55677 into the next 7.2.15.x maintenance release
- HOTREQ-1468: keytab does not exist error - ENGESC-20539

Known issue

OPSAPS-68584

NiFi Registry fails with Postgres 14 DB due to unable to obtain connection from Flyway DataSource.

You can either use older version of Postgres (Postgres 11) or update the postgresql jdbc driver under /usr/share/java on the machine where nifi-registry is installed (the machine which has CM server) and ensure that the driver file access right is approriate.

For example,

```
wget -O /usr/share/java/postgresql-jdbc.jar https://jdbc.postgre
sql.org/download/postgresql-42.5.4.jar
```

CDPD-61655

Hue - You can not create file/folder with non-ascii char on Azure.

None.

Fixed Issues In Cloudera Runtime 7.2.15.400

You can review the list of reported issues and their fixes in Cloudera Runtime 7.2.15.400.

CDH

- HOTREQ-1590 Need hotfix for HIVE-21100 on CDP 7.2.15.0

CFM

- HOTREQ-1587 Ship NIFI-12160 for 7.2.17, 7.2.16, and 7.2.15

Known issue

OPSAPS-68584

NiFi Registry fails with Postgres 14 DB due to unable to obtain connection from Flyway DataSource.

You can either use older version of Postgres (Postgres 11) or update the postgresql jdbc driver under /usr/share/java on the machine where nifi-registry is installed (the machine which has CM server) and ensure that the driver file access right is approriate.

For example,

```
wget -O /usr/share/java/postgresql-jdbc.jar https://jdbc.postgre
sql.org/download/postgresql-42.5.4.jar
```

CDPD-61655

Hue - You can not create file/folder with non-ascii char on Azure.

None.

Service Pack in Cloudera Runtime 7.2.15

You can review the list of CDP Public Cloud hotfixes rolled into Cloudera Runtime 7.2.15. 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-5132

Known Issues In Cloudera Runtime 7.2.15

You must be aware of the known issues and limitations, the areas of impact, and workaround in Cloudera Runtime 7.2.15.

Known Issues in Apache Atlas

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

CDPD-35691:Affected version - 7.2.15: Task API intermittently returns a blank json object, even when there are no pending tasks.

None

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 blog 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 Cruise Control

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

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

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

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

Configure Cruise Control to use to use the Cruise Control metrics reporter (default). This issue is not present if this metric reporter is used.

1. In Cloudera Manager, select the Cruise Control service.
2. Go to Configuration.
3. Find the Metric Reporter property.
4. Select the Cruise Control metrics reporter option.
5. Restart the Cruise Control service.

OPSAPS-68148: Cruise Control rack aware goal upgrade handler

The goal sets in Cruise Control, which include the default, supported, hard, self-healing and anomaly detection goals, might be overridden to their default value after a cluster upgrade if the goals have been customized.

Create a copy from the values of the goal lists before upgrading your cluster, and add the copied values to the goal lists after upgrading the cluster. Furthermore, you must rename any mentioning of com.linkedin.kafka.cruisecontrol.analyzer.goals.RackAwareGoal to com.linkedin.kafka.cruisecontrol.analyzer.goals.RackAwareDistributionGoal as Cruise Control will not be able to start otherwise.

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.

- CDPD-49281: DAS WebApp logs are not captured in the var/logs/das/ directory, as expected.

  Workaround: To obtain the DAS WebApp logs, check the stderr.log file in the runtime process directory for the DAS WebApp.
- You may not be able to add or delete columns or change the table schema after creating a new table using the upload table feature.
- For clusters secured using Knox, you see the HTTP 401: Forbidden error message when you click the DAS quick link from Cloudera Manager and are unable to log into DAS.

  Workaround: The admin user will need to provide the DAS URL from the Knox proxy topology to the users needing access to DAS.
- The download logs feature may not return the YARN application logs on a Kerberized cluster. When you download the logs, the logs contain an error-reports.json file which states that no valid Kerberos tokens are available.

  Workaround: An admin user with access to the machine can use the kinit command as a hive user with hive service user keytabs and trigger the download.
- The task logs for a particular task may not be available in the task swimlane. And the zip file generated by download logs artifact may not have task logs, but instead contain an error-reports.json file with the error log of the download failures.

- You may not see any data for a report for any new queries that you run. This can happen especially for the last one day's report.

  Workaround:

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

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

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

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

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

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

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

Known Issues in Apache HBase

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

OpDB Data Hub cluster fails to initialize if you are reusing a cloud storage location that was used by an older OpDB Data Hub cluster. You cannot create a new OpDB Data Hub cluster because OpDB Data Hub template is not suported.

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

When using Replication Manager for Apache HBase (HBase) snapshot replication, data loss will occur if both of the following conditions are met: (i) the external account used for the operation has delete access to the target storage location, and (ii) the snapshot export fails. If these conditions are met, the cleanup operation, which is automatically performed after the failure, would delete all data in the root folder of the snapshot, not only the snapshot files. If the user account does not have the delete permission on the target folder, the data remains unaffected.

Knowledge article

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

Known Issues in HDFS

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

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

None.

Unsupported Features

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

- ACLs for the NFS gateway (HADOOP-11004)
- Aliyun Cloud Connector (HADOOP-12756)
- Allow HDFS block replicas to be provided by an external storage system (HDFS-9806)
- Consistent standby Serving reads (HDFS-12943)
- Cost-Based RPC FairCallQueue (HDFS-14403)

- HDFS Router Based Federation (HDFS-10467)
- More than two NameNodes (HDFS-6440)
- NameNode Federation (HDFS-1052)
- NameNode Port-based Selective Encryption (HDFS-13541)
- Non-Volatile Storage Class Memory (SCM) in HDFS Cache Directives (HDFS-13762)
- OpenStack Swift (HADOOP-8545)
- SFTP FileSystem (HADOOP-5732)
- Storage policy satisfier (HDFS-10285)

Technical Service Bulletins

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

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

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

Knowledge article

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

Known Issues in Apache Hive

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

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

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

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

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

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

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

HiveServer Web UI displays incorrect data

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

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

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

Use one of the workarounds:

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

Technical Service Bulletins

TSB 2022-567: Potential Data Loss due to CTLT HBaseStorageHandler failure dropping underlying HBase table while rollback

If the create table target_table like source table command (CTLT) fails and the source table is HBaseStorageHandler-based table, the HBaseMetaHook rollback logic deletes the underlying HBase table, resulting in potential data loss.

Upstream JIRA

HIVE-25989

Knowledge article

For the latest update on this issue, see the corresponding Knowledge article: TSB 2022-567: Potential Data Loss due to CTLT HBaseStorageHandler failure dropping underlying HBase table while rollback

TSB 2022-600: Renaming translated external partition table shows empty records in Apache Hive

If an Apache Hive partitioned table is renamed, it can cause data loss due to the location being incorrectly translated at the Hive Metastore (HMS) translation layer in the legacy config mode.

Scenario:

- The following configurations are set:
  - hive.create.as.external.legacy=true
  - hive.created.as.acid=true
- The following processes are executed:
  - Creation of new partition table
  - Data is loaded on new table
  - Table is renamed
  - Scan/view after rename of the same table returns empty records

Example:

- The following kind of query is affected:

```
CREATE TABLE foo (i1 int) PARTITIONED BY (i2 string);
                  INSERT INTO foo VALUES (1,’foo’);
                  ALTER TABLE foo RENAME TO foo_renamed;
                  SELECT  * FROM foo_renamed; //returns empty
 records
```

- The following kind of query is not affected:

```
CREATE EXTERNAL foo (i1 int) PARTITIONED BY (i2 string);
                  INSERT INTO foo VALUES (1,’foo’);
                  ALTER TABLE foo RENAME TO foo_renamed;
                  SELECT  * FROM foo_renamed; //returns 1 reco
rd
```

Upstream JIRA

HIVE-26158

Knowledge article

For the latest update on this issue, see the corresponding Knowledge article: TSB 2022-600: Renaming translated external partition table shows empty records in Apache Hive

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

TSB 2023-653: Cleaner causes data loss when processing an aborted dynamic partitioning transaction

If the compaction-cleaner is enabled, 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 2023-653: Cleaner causes data loss when processing an aborted dynamic partitioning transaction

Known Issues in Hue

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

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

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

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

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

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

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

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

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

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

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

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

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

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

3. Back up the hue directory:

   ```
   cp -R hue hue_orginal
   ```

4. Change to the hue directory:

   ```
   cd hue
   ```

5. Install the Wheel package using pip:

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

   The Wheel package is used to avoid recompiling your software during every install.
6. Install the Python Setuptools package for Hue as follows:

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

7. Install Tablib version 0.12.1 as follows:

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

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

Impala SELECT table query fails with UTF-8 codec error

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

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

Hue Importer is not supported in the Data Engineering template

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

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

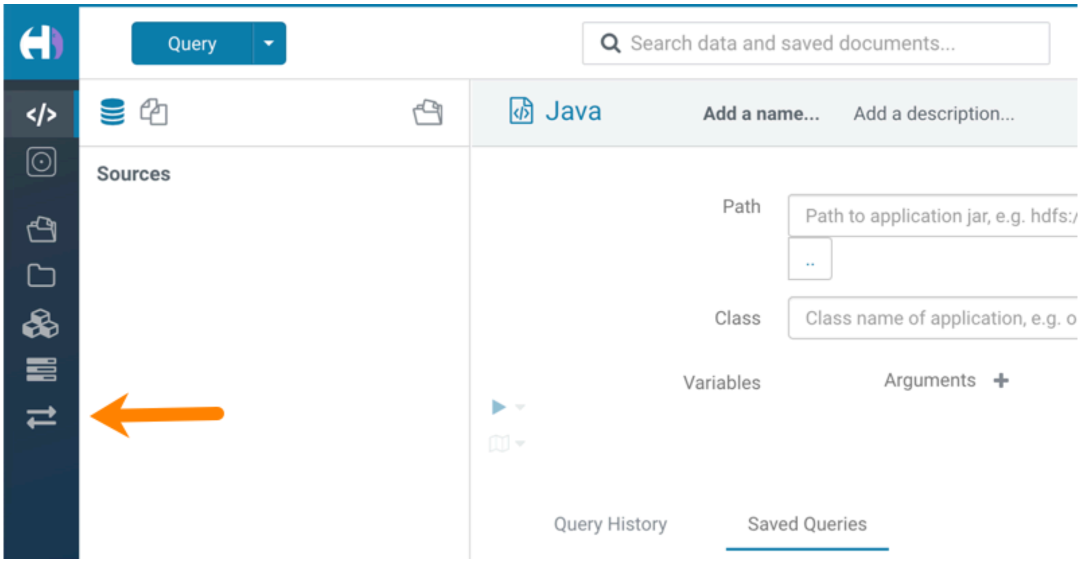

Hue Load Balancer role fails to start after upgrade to Cloudera Runtime 7 or you get the "BalancerMember worker hostname too long" error

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

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

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

```
com.sequenceiq.cloudbreak.cm.ClouderaManagerOperationFailedExcep
tion: Cluster template install failed: [Command [Start], with id
 [1234567890] failed:
Failed to start role., Command [Start], with id [1234567890] fail
ed: 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.confi
g.ConfigGenException: Unable to generate config file hue.ini
```

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

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

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

For example, cdp-123456-scale.

Unsupported features

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

You can export Oozie workflows, schedules, and bundles from Hue and import them only within the same cluster if the cluster is unchanged. You can migrate bundle and coordinator jobs with their workflows only if their arguments have not changed between the old and the new cluster. For example, hostnames, NameNode, Resource Manager names, YARN queue names, and all the other parameters defined in the workflow.xml and job.properties files.

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

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

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

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

None.

Technical Service Bulletins

TSB 2024-723: Hue RAZ is using logger role to Read and Upload/Delete (write) files

When using Cloudera Data Hub for Public Cloud (Data Hub) on Amazon Web Services (AWS), users can use the Hue File Browser feature to access the filesystem, and if permitted, read and write directly to the related S3 buckets. As AWS does not provide fine-grained access control, Cloudera Data Platform administrators can use the Ranger Authorization Service (RAZ) capability to take the S3 filesystem, and overlay it with user and group specific permissions, making it easier to allow certain users to have limited permissions, without having to grant those users permissions to the entire S3 bucket.

This bulletin describes an issue when using RAZ with Data Hub, and attempting to use fine-grained access control to allow certain users write permissions.

Through RAZ, an administrator may, for a particular user, specify permissions more limited than what AWS provides for an S3 bucket, allowing the user to have read/write (or other similar fine grained access) permissions on only a subset of the files and directories within that bucket. However, under specific conditions, it is possible for such user to be able to read and write to the entire S3 bucket through Hue, due to Hue using the logger role (which will have full read/write to the S3 bucket) when using Data Hub with a RAZ enabled cluster. This problem also can affect the Hue service itself, by affecting proper access to home directories causing the service role to not start.

The root cause of this issue is, when accessing Amazon cloud resources, Hue uses the AWS Boto SDK library. This AWS Boto library has a bug that restricts permissions in certain AWS regions in such a way that it provides access to users who should not have it, regardless of RAZ settings. This issue only affects users in specific AWS regions, listed below, and it does not affect all AWS customers.

Knowledge article

For the latest update on this issue see the corresponding Knowledge Article: TSB 2024-723: Hue Raz is using logger role to Read and Upload/Delete (write) files.

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.

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

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

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

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

This vulnerability is fixed upstream under IMPALA-10600

Knowledge article

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

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.

Known Issues in Apache Kafka

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

Known Issues

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

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

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

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

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

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

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

Certain Kafka command line tools require direct access to Zookeeper

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

- kafka-reassign-partitions

None

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

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

None

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

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

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

KAFKA-2561: Performance degradation when SSL Is enabled

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

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

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

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

None

OPSAPS-63640: Monitoring a high number of Kafka producers might cause Cloudera Manager to slow down and run out of memory

This issue has two workarounds. You can either configure a Kafka producer metric allow list or completely disable producer metrics.

- Configure a Kafka producer metric allow list:

  A producer metric allow list can be configured by adding the following properties to Kafka Broker Advanced Configuration Snippet (Safety Valve) for kafka.properties.

  ```
  producer.metrics.whitelist.enabled=true
  producer.metrics.whitelist=[***ALLOW LIST REGEX***]
  ```

  Replace *[***ALLOW LIST REGEX***]* with a regular expression matching the client.id of the producers that you want to add to the allow list. This regular expression uses the java.util.regex. Pattern class to compile the regular expression, and uses the match() method on the client.id to determine whether it fits the regular expression.

  Once configured, the metrics of producers whose client.id does not match the regular expression provided in producer.metrics.whitelist are filtered.Kafka no longer reports these metrics through the HTTP metrics endpoint. Additionally, existing metrics of the producers whose client.id does not match the regular expression are deleted.

  Because the allow list filters metrics based on the client.id of the producers, you must ensure that the client.id property is specified in each producer's configuration. Automatically generated client IDs might cause the number of unnecessary metrics to increase even if an allow list is configured.
- Completely disable producer metrics:

  Producer metrics can be completely disabled by unchecking the Enable Producer Metrics Kafka service property.

CDPD-39354: Kafka Connect connectors and tasks fail to start

Kafka Connect initializes the Secrets Storage late, causing some connectors and tasks to fail at startup. This is a timing issue, and does not occur deterministically.

Restart the failed connectors and tasks manually, or do not use the Secrets Storage feature in the configurations.

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

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

None

CDPD-45958: Kafka client JAAS override policy validation is incorrect

The JAAS override filter policy refuses configurations if the configuration contains an unknown field instead of only refusing based on known fields with invalid values.

None

CDPD-48822: AvroConverter ignores default values when converting from Avro to Connect schema

AvroConverter does not propagate field default values when converting Avro schemas to Connect schemas.

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.
- The Kafka default authorizer is not supported. This includes setting ACLs and all related APIs, broker functionality, and command-line tools.
- SASL/SCRAM is only supported for delegation token based authentication. It is not supported as a standalone authentication mechanism.
- The Cloudera `AvroConverter` (`com.cloudera.dim.kafka.connect.converts.AvroConverter`) is not supported with the Debezium Kafka Connect connectors shipped in Cloudera Runtime.

Limitations

Collection of Partition Level Metrics May Cause Cloudera Manager's Performance to Degrade

If the Kafka service operates with a large number of partitions, collection of partition level metrics may cause Cloudera Manager's performance to degrade.

If you are observing performance degradation and your cluster is operating with a high number of partitions, you can choose to disable the collection of partition level metrics.

⚠ **Important:** If you are using SMM to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both SMM and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, SMM will not be able to display information about partitions. In addition, Cruise Control will not operate properly.

Complete the following steps to turn off the collection of partition level metrics:

1. Obtain the Kafka service name:

   a. In Cloudera Manager, Select the Kafka service.
   b. Select any available chart, and select Open in Chart Builder from the configuration icon drop-down.
   c. Find $SERVICENAME= near the top of the display.

   The Kafka service name is the value of $SERVICENAME.

2. Turn off the collection of partition level metrics:

    a. Go to HostsHosts Configuration.
    b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

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

        ```
        [KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_ent
        ity_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.

Technical Service Bulletins

TSB 2022-614: Kafka policy "user auto-creation" does not work in Ranger in CDP Public Cloud 7.2.15

Creating a completely new Cloudera Data Platform (CDP) Public Cloud 7.2.15 Streams Messaging Light Duty Data Hub cluster fails after the Data Lake upgrade to 7.2.15. Note that the Data Hub cluster is created, but it looks unusable because of the lack of permissions.

New user principals are added to Apache Kafka (Kafka) policies (cc_metric_reporter and kafka_mirror_maker) in CDP Public Cloud version 7.2.14 as part of new features. Whenever a new Data Hub cluster is installed, its Kafka service is started for the first time, it will try to create all the default Kafka related policies automatically. If any of the users referred to in the policies does not exist in Apache Ranger (Ranger), it will refuse creating any of the policies in CDP Public Cloud version 7.2.15 for the new Data Hub cluster and the cluster will look unusable.

This is because from CDP Public Cloud 7.2.15 onwards, Ranger only lets administrators create new users. Automatic user creation works in CDP Public Cloud 7.2.14, so the affected customers will depend on which versions of Streams Messaging Data Hub clusters and Data Lakes they used earlier and how they used them.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: TSB 2022-614: Kafka policy "user auto-creation" does not work in Ranger in CDP Public Cloud 7.2.15

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.

Technical Service Bulletins

TSB 2023-630: Apache Knox - Server-side Request Forgery in host parameter

When authenticated to an Apache Knox (Knox) protected endpoint, such as Apache HBase (HBase), modifying the host parameter by adding an external host causes Knox to unexpectedly send a request to the external host which includes the user's cookies. A malicious actor may present this request URL to the user through an XSS attack or phishing campaign.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: TSB 2023-630: Apache Knox - Server-side Request Forgery in host parameter

Known Issues in Apache Kudu

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

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

None

Kudu HMS Sync is disabled and is not yet supported

None

Known Issues in Apache Oozie

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

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

None

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

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

Use Spark action in cluster mode.

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

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

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

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

   Set the following if you are using Amazon S3:

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

   Set the following if you are using ABFS:

   ```
   <property>
         <name>oozie.service.HadoopAccessorService.fs.abfs</name>
         <value>fs.azure.buffer.dir=/tmp/abfs</value>
      </property>
      <property>
         <name>oozie.service.HadoopAccessorService.fs.abfss</name
   >
         <value>fs.azure.buffer.dir=/tmp/abfss</value>
      </property>
   ```
4. Enter a Reason for change, and then click Save Change to commit the changes.
5. Restart the Oozie service.

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

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

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

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

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

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

Unsupported Feature

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

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

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

BUG-123856: Upgrade fails while configuring Oozie server.

None

Known Issues in Apache Phoenix

There are no known issues for Phoenix in Cloudera Runtime 7.2.15.

Known Issues in Apache Ranger

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

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

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

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

1. On the Configuration tab in the applicable service, select Advanced under CATEGORY.
2. Click the + icon for the <service_name> Advanced Configuration Snippet (Safety Valve) for ranger-<service_name>-audit.xml property.
3. Enter the following property in the Name box:

   xasecure.audit.destination.hdfs.file.rollover.sec.
4. Enter the desired rollover interval (in seconds) in the Value box. For example, if you specify 180, the audit log data is updated every 3 minutes.
5. Click Save Changes and restart the service.

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

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

Use only unique key names when creating keys.

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

There are two possible workarounds to resolve this issue:

1. Update database entries (Recommended):

   - select * from x_ranger_global_state where state_name='RangerRole';
   - update x_ranger_global_state set app_data='{"Version":"2"}' where state_name='RangerRole';

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

Known Issues in Schema Registry

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

CDPD-54379: KafkaJsonSerializer and KafkaJsonDeserializer do not allow null values

`KafkaJsonSerializer` and `KafkaJsonDeserializer` do not allow the data to be null, resulting in a `NullPointerException` (NPE).

None.

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

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

Restart Schema Registry after group membership changes.

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

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

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

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

CDPD-58265: Schema Registry Client incorrectly applies SSL configuration

The Cloudera distributed Schema Registry Java client might fail to apply the SSL configurations correctly with concurrent access in Jersey clients due to a Jersey issue related to JDK.

Before using HttpsURLConnection in any form concurrently, call javax.net.ssl.HttpsURLConnection.getDefaultSSLSocketFactory() once in the custom client application.

CDPD-48822: AvroConverter ignores default values when converting from Avro to Connect schema

AvroConverter does not propagate field default values when converting Avro schemas to Connect schemas.

None

CDPD-55381: Schema Registry issues authentication cookie for the authorized user, not for the authenticated one

When the authenticated user is different from the authorized user, which can happen when Schema Registry is used behind Knox, authorization issues can occur for subsequent requests as the authentication cookie in Schema Registry stores the authorized user.

Access Schema Registry directly, without using Knox, if possible. If not, ensure that the name of the end user that tries to connect does not begin with knox.

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

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

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

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

None.

CDPD-48853: Schemas created with the Confluent Schema Registry API cannot be viewed in the UI

Schemas created in Cloudera Schema Registry using the Confluent Schema Registry API are not visible in the Cloudera Schema Registry UI.

In addition, the `/api/v1/schemaregistry/search/schemas/aggregated` endpoint of the Cloudera Schema Registry API does not return schemas created with the Confluent Schema Registry API.

A typical case where this issue can manifest is when you are using the Confluent Avro converter for SerDes in a Kafka Connect connector and the connector connects to Cloudera Schema Registry. That is, the key.converter and/or value.converter properties of the connector are set to `io.confluent.connect.avro.AvroConverter`, and key.converter.schema.registry.url and/or value.converter.schema.registry.url are set to a Cloudera Schema Registry server URL.

None.

CDPD-58949: Schemas are de-duplicated on import

On import, Schema Registry de-duplicates schema versions based on their fingerprints. This means that schemas which are considered functionally equivalent in SR get de-duplicated. As a result, some schema versions are not created, and their IDs do not become valid IDs in SR.

None.

CDPD-58990: getSortedSchemaVersions method orders by schemaVersionId instead of version number

On validation, Schema Registry orders schema versions based on ID instead of version number. In some situations, this can cause validation with the LATEST level to compare the new schema version to a non-latest version.

This situation can occur when an older version of a schema has a higher ID than the newer version of a schema, for example, when the older version is imported with an explicit ID.

None.

Known Issues in Cloudera Search

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

Note: From Cloudera Runtime 7.2.15 and higher, Cloudera Search known issues are listed under Apache Solr.

Known Issues in Apache Solr

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

Known Issues

Changing the default value of Client Connection Registry HBase configuration parameter causes HBase MRIT job to fail

If the value of the HBase configuration property Client Connection Registry is changed from the default ZooKeeper Quorum to Master Registry then the Yarn job started by HBase MRIT fails with a similar error message:

```
Caused by: org.apache.hadoop.hbase.exceptions.MasterRegistryFetc
hException: Exception making rpc to masters [quasar-bmyccr-2.qua
sar-bmyccr.root.hwx.site,22001,-1]
        at org.apache.hadoop.hbase.client.MasterRegistry.lambda$g
roupCall$1(MasterRegistry.java:244)
        at org.apache.hadoop.hbase.util.FutureUtils.lambda$addLi
stener$0(FutureUtils.java:68)
        at java.util.concurrent.CompletableFuture.uniWhenCompl
ete(CompletableFuture.java:774)
        at java.util.concurrent.CompletableFuture.uniWhenComplet
eStage(CompletableFuture.java:792)
        at java.util.concurrent.CompletableFuture.whenComplete(Co
mpletableFuture.java:2153)
        at org.apache.hadoop.hbase.util.FutureUtils.addListener(F
utureUtils.java:61)
        at org.apache.hadoop.hbase.client.MasterRegistry.groupCa
ll(MasterRegistry.java:228)
        at org.apache.hadoop.hbase.client.MasterRegistry.call(Ma
sterRegistry.java:265)
        at org.apache.hadoop.hbase.client.MasterRegistry.getMetaR
egionLocations(MasterRegistry.java:282)
        at org.apache.hadoop.hbase.client.ConnectionImplementati
on.locateMeta(ConnectionImplementation.java:900)
        at org.apache.hadoop.hbase.client.ConnectionImplementat
ion.locateRegion(ConnectionImplementation.java:867)
        at org.apache.hadoop.hbase.client.ConnectionImplementati
on.relocateRegion(ConnectionImplementation.java:850)
        at org.apache.hadoop.hbase.client.ConnectionImplementat
ion.locateRegionInMeta(ConnectionImplementation.java:981)
        at org.apache.hadoop.hbase.client.ConnectionImplementa
tion.locateRegion(ConnectionImplementation.java:870)
        at org.apache.hadoop.hbase.client.RpcRetryingCallerWith
ReadReplicas.getRegionLocations(RpcRetryingCallerWithReadReplica
s.java:319)
        ... 21 more
Caused by: org.apache.hadoop.hbase.client.RetriesExhaustedExcept
ion: Failed contacting masters after 1 attempts.
Exceptions:
java.io.IOException: Call to address=quasar-bmyccr-2.quasar-bmy
ccr.root.hwx.site/172.27.19.4:22001 failed on local exception: j
```

```
ava.io.IOException: java.lang.RuntimeException: Found no valid a
uthentication method from options
        at org.apache.hadoop.hbase.client.MasterRegistry.lambda
$groupCall$1(MasterRegistry.java:243)
        ... 35 more
```

Add the following line to the MRIT command line:

```
-D 'hbase.client.registry.impl=org.apache.hadoop.hbase.client.ZK
ConnectionRegistry'
```

Apache Tika upgrade may break morphlines indexing

The upgrade of Apache Tika from 1.27 to 2.3.0 brought potentially breaking changes for morphlines indexing. Duplicate/triplicate keys names were removed and certain parser class names were changed (For example, org.apache.tika.parser.jpeg.JpegParser changed to org.apache.tika.parser.image.JpegParser).

To avoid morphline commands failing after the upgrade, do the following:

- Check if key name changes affect your morphlines. For more information, see *Removed duplicate/triplicate keys* in Migrating to Tika 2.0.0.
- Check if the name of any parser you use has changed. For more information, see the Apache Tika API documentation.

Update your morphlines if necessary.

CDPD-28432: HBase Lily indexer REST port does not support SSL

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

Switch off the REST API setting the hbaseindexer.httpserver.disabled environment parameter to true (by default this is false). This switches off the REST interface, so noone can use the --http argument when using the hbase-indexer command line tool. This also means that users need to authenticate as hbase user in order to use the hbase-indexer tool.

CDH-77598: Indexing fails with socketTimeout

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

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

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

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

CDPD-20577: Splitshard operation on HDFS index checks local filesystem and fails

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

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

- On nonsecure clusters:

  ```
  curl 'http://$[***SOLR_SERVER_HOSTNAME***]:8983/so
  lr/admin/collections?action=SPLITSHARD&collectio
  ```

```
n=[***COLLECTION_NAME***]&shard=[***SHARD_TO_SPLIT***]&skipFre
eSpaceCheck=true'
```

- On secure clusters:

```
curl -k -u : --negotiate 'http://
$[***SOLR_SERVER_HOSTNAME***]:8985/solr/admin/collections
?action=SPLITSHARD&collection=[***COLLECTION_NAME***]&sha
rd=[***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 succesfully, 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.

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

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

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

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

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

CDH-22190: CrunchIndexerTool which includes Spark indexer requires specific input file format specifications

If the --input-file-format option is specified with CrunchIndexerTool, then its argument must be text, avro, or avroParquet, rather than a fully qualified class name.

None

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

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

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

CDH-26856: Field value class guessing and Automatic schema field addition are not supported with the MapReduceIndexerTool nor with the HBaseMapReduceIndexerTool.

The MapReduceIndexerTool and the HBaseMapReduceIndexerTool can be used with a Managed Schema created via NRT indexing of documents or via the Solr Schema API. However, neither tool supports adding fields automatically to the schema during ingest.

Define the schema before running the MapReduceIndexerTool or HBaseMapReduceIndexerTool. In non-schemaless mode, define in the schema using the schema.xml file. In schemaless mode, either define the schema using the Solr Schema API or index sample documents using NRT indexing before invoking the tools. In either case, Cloudera recommends that you verify that the schema is what you expect, using the List Fields API command.

CDH-19407: The Browse and Spell Request Handlers are not enabled in schemaless mode

The Browse and Spell Request Handlers require certain fields to be present in the schema. Since those fields cannot be guaranteed to exist in a Schemaless setup, the Browse and Spell Request Handlers are not enabled by default.

If you require the Browse and Spell Request Handlers, add them to the solrconfig.xml configuration file. Generate a non-schemaless configuration to see the usual settings and modify the required fields to fit your schema.

CDH-17978: Enabling blockcache writing may result in unusable indexes.

It is possible to create indexes with solr.hdfs.blockcache.write.enabled set to true. Such indexes may appear corrupt to readers, and reading these indexes may irrecoverably corrupt indexes. Blockcache writing is disabled by default.

None

CDH-58276: Users with insufficient Solr permissions may receive a "Page Loading" message from the Solr Web Admin UI.

Users who are not authorized to use the Solr Admin UI are not given a page explaining that access is denied to them, instead receive a web page that never finishes loading.

None

CDH-15441: Using MapReduceIndexerTool or HBaseMapReduceIndexerTool multiple times may produce duplicate entries in a collection.

Repeatedly running the MapReduceIndexerTool on the same set of input files can result in duplicate entries in the Solr collection. This occurs because the tool can only insert documents and cannot update or delete existing Solr documents. This issue does not apply to the HBaseMapReduceIndexerTool unless it is run with more than zero reducers.

To avoid this issue, use HBaseMapReduceIndexerTool with zero reducers. This must be done without Kerberos.

CDH-58694: Deleting collections might fail if hosts are unavailable.

It is possible to delete a collection when hosts that host some of the collection are unavailable. After such a deletion, if the previously unavailable hosts are brought back online, the deleted collection may be restored.

Ensure all hosts are online before deleting collections.

Unsupported Features

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

- Package Management System
- HTTP/2
- Solr SQL/JDBC
- Graph Traversal
- Cross Data Center Replication (CDCR)
- SolrCloud Autoscaling
- HDFS Federation
- Saving search results
- Solr contrib modules (Spark, MapReduce and Lily HBase indexers are not contrib modules but part of the Cloudera Search product itself, therefore they are supported).

Known Issues in Apache Spark

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

CDPD-217: The Apache Spark connector is not supported

The old *Apache Spark - Apache HBase Connector* (shc) is not supported in CDP releases.

Use the new HBase-Spark connector shipped in CDP release.

CDPD-3038: Launching pyspark displays several HiveConf warning messages

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

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

These errors can be safely ignored.

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

Use a different compression algorithm.

CDPD-3783: Cannot create databases from Spark

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

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

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

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

Known Issues for Apache Sqoop

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

Using direct mode causes problems

Using direct mode has several drawbacks:

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

Stop using direct mode. Do not use the --direct option in Sqoop import or export commands.

CDPD-3089: Avro, S3, and HCat do not work together properly

Importing an Avro file into S3 with HCat fails with Delegation Token not available.

Parquet columns inadvertently renamed

Column names that start with a number are renamed when you use the --as-parquetfile option to import data.

Prepend column names in Parquet tables with one or more letters or underscore characters.

Importing Parquet files might cause out-of-memory (OOM) errors

Importing multiple megabytes per row before initial-page-run check (ColumnWriter) can cause OOM. Also, rows that vary significantly by size so that the next-page-size check is based on small rows, and is set very high, followed by many large rows can also cause OOM.

None

Known issues in Streams Messaging Manager

Learn about the known issues for Streams Messaging Manager in Cloudera Runtime 7.2.15.

CDPD-39826: The Restart button for the ConnectorTasks is permanently disabled

On the ConnectorDetails page the Restart button for the tasks within the connector is permanently disabled.

Restart the whole Connector.

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

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

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

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

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

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

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

View the SMM UI logs on the host.

CDPD-46728: SMM UI shows the consumerGroup instead of the instances on the Profile page's right hand side

On the ConsumerGroupDetail page, SMM UI shows the group instead of its instances on the right hand side table.

None.

Limitations

CDPD-36422: 1MB flow.snapshot freezes Safari

While importing large connector configurations, flow.snapshots reduces the usability of the Streams Messaging Manager when using Safari browser.

Use a different browser (Chrome/Firefox/Edge).

Known Issues in Streams Replication Manager

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

Known Issues

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

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

None

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

If a topic was originally replicated but was later 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, it's Partitions and Consumer Groups will be 0, its Throughput, Replication Latency and Checkpoint Latency will show N/A.

None

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

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

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

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

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

OPSAPS-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

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

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

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

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

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

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

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

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

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

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

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

CDPD-60426: Configuration changes are lost following a rolling restart of the service

In certain cases, SRM might fail to apply configuration updates if the service is restarted with a rolling restart. In a case like this, configuration changes are ignored without any warning or indication. This issue also affects rolling upgrades.

When restarting the service, use Actions Restart instead of Actions Rolling Restart after making configuration changes. When upgrading a cluster, ensure that SRM is not restarted with a rolling restart.

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 *[***CONFIG LEVEL PREFIX***]*.isolation.level=read_committed to the Streams Replication Manager's Replication Configs SRM service property in Cloudera Manger. The isolation.level property can be set on a global connector or replication level. For example:

```
#Global connector level
connectors.consumer.isolation.level=read_committed
#Replication level
uswest->useast.consumer.isolation.level=read_committed
```

SRM checkpointing is not supported for transactional source topics

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

Known Issues in MapReduce and YARN

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

Known Issues

YARN cannot start if Kerberos principal name is changed

If the Kerberos principal name is changed in Cloudera Manager after launch, YARN will not be able to start. In such case the keytabs can be correctly generated but YARN cannot access ZooKeeper with the new Kerberos principal name and old ACLs.

There are two possible workarounds:

- Delete the znode and restart the YARN service.
- Use the reset ZK ACLs command. This also sets the znodes below /rmstore/ZKRMStateRoot to world:anyone:cdrwa which is less secure.

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

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

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

JobHistory URL mismatch after server relocation

After moving the JobHistory Server to a new host, the URLs listed for the JobHistory Server on the ResourceManager web UI still point to the old JobHistory Server. This affects existing jobs only. New jobs started after the move are not affected.

For any existing jobs that have the incorrect JobHistory Server URL, there is no option other than to allow the jobs to roll off the history over time. For new jobs, make sure that all clients have the updated mapred-site.xml that references the correct JobHistory Server.

CDH-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.

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.

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

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-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-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.

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

In the current implemention, 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-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-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-7586: Max Parallel Apps cannot be changed for root queue

The Queue Manager UI may show incorrect value for the Maximum Parallel Applications property for the root queue. The value can be changed with the QM UI, but the UI will always show 2147483647.

Check the Maximum Parallel Applications property for the root queue manually:

1. In Cloudera Manager, select to the YARN service.
2. Click Configuration.
3. Find the Capacity Scheduler Configuration Advanced Configuration Snippet (Safety Valve) property.
4. Find the yarn.scheduler.capacity.root.max-parallel-apps property in the safety valve.

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.

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.

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.

COMPX-8687: Missing access check for getAppAttemps

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

Technical Service Bulletins

TSB 2023-641: InvalidClassException while editing queue configurations in YARN Queue Manager UI

Under situations described below, a user may encounter the following error message while editing queue configurations in the Apache Hadoop YARN (YARN) Queue Manager UI:

Failed to update queue configuration

Modify queue operation failed. Queue configuration in Cloudera Manager is inconsistent with YARN. Try restarting YARN.

Knowledge article

For the latest update on this issue see the corresponding Knowledge article: TSB 2022-641: InvalidClassException while editing queue configurations in YARN Queue Manager UI.

Unsupported Features

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

- Application Timeline Server (ATSv2 and ATSv1)
- Container Resizing
- Distributed or Centralized Allocation of Opportunistic Containers
- Distributed Scheduling
- Docker on YARN (DockerContainerExecutor) on Data Hub clusters
- Fair Scheduler
- GPU support for Docker
- Hadoop Pipes
- Native Services
- Pluggable Scheduler Configuration
- Queue Priority Support
- Reservation REST APIs
- Resource Estimator Service
- Resource Profiles
- (non-Zookeeper) ResourceManager State Store
- Rolling Log Aggregation
- Shared Cache
- YARN Federation
- Moving jobs between queues

Known Issues in Apache Zeppelin

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

CDPD-3090: Due to a configuration typo, functionality involving notebook repositories does not work

Due to a missing closing brace, access to the notebook repositories API is blocked by default.

From the CDP Management Console, go to Cloudera Manager for the cluster running Zeppelin. On the Zeppelin configuration page (Zeppelin serviceConfiguration), enter shiro urls in the Search field, and then add the missing closing brace to the notebook-repositories URL, as follows:

```
/api/notebook-repositories/** = authc, roles[{{zeppelin_admin_gr
oup}}]
```

Click Save Changes, and restart the Zeppelin service.

CDPD-2406: Logout button does not work

Clicking the Logout button in the Zeppelin UI logs you out, but then immediately logs you back in using SSO.

Close the browser.

Known Issues in Apache ZooKeeper

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

Zookeeper-client does not use ZooKeeper TLS/SSL automatically

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

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

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

Behavioral Changes In Cloudera Runtime 7.2.15

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 7.2.15.

Behavioral Changes in Apache Kafka

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

Summary:

JMX authentication is enabled by default for the Kafka service. Additionally random passwords are now generated for both JMX users by default.

Previous behavior:

JMX authentication (Enable Authenticated Communication with the JMX Agent) was disabled by default.

New behavior:

JMX authentication (Enable Authenticated Communication with the JMX Agent) is now enabled by default. By default random passwords are generated for both JMX users.

Deprecation Notices In Cloudera Runtime 7.2.15

Certain features and functionalities have been removed or deprecated in Cloudera Runtime 7.2.15. 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 7.2.15. You must review these changes along with the information about the features in Kudu that will be removed or deprecated in a future release.

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

Deprecation Notices for Apache Kafka

Certain features and functionality in Apache Kafka are deprecated or removed in Cloudera Runtime 7.2.15. 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-re assign-partitions. Cloudera recommends that you use the --bootstrap-server option instead.

Deprecation notices in Streams Messaging Manager

Certain features and functionalities in Streams Messaging Manager (SMM) are deprecated or removed in Cloudera Runtime 7.2.15. You must review these changes along with the information about the features in SMM that will be removed or deprecated in a future release.

The lineage and aggregated metrics endpoints have their versions changed.

- The experimental v1 lineage endpoints changes include:
    - /api/v1/admin/lineage/producers/{producerId} is changed to /api/v2/admin/lineage/producers/{producerId}. The response object is not backward compatible, however, all the information is still present in a better structure.
    - /api/v1/admin/lineage/consumersGroup/{consumersGroupId} is changed to /api/v2/admin/lineage/consumersGroup/{consumersGroupId}. The response object is not backward compatible, however, all the information is still present in a better structure.
    - /api/v1/admin/lineage/partitions/{topicName} is changed to /api/v2/admin/lineage/partitions/{topicName}. The response object is unchanged.
    - /api/v1/admin/lineage/partitions/{topicName}/{partitionId} is changed to /api/v2/admin/lineage/partitions/{topicName}/{partitionId}. The response object is unchanged.
- The aggregated metrics endpoints changes include:
    - Every endpoint within /api/v1/admin/metrics/aggregated/* went through a version change to /api/v2/admin/metrics/aggregated/*.
    - The old endpoints are not accessible any more.
    - Most notably the partitionMetrics field is removed from all the endpoints and the information within is not accessible any more.