Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 Release Notes

Date published: 2020-07-28
Date modified: 2026-08-28



https://docs.cloudera.com/

Legal Notice

© Cloudera Inc. 2026. All rights reserved.

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

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

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

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

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

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

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

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

Contents

Overview 7

What's new in Platform Support 8

Runtime API Compatibility 9
- Introduction to the API compatibility 9
- Client-server compatibility 9

What's New in Cloudera Runtime 7.3.2 11
- Atlas 11
- What's New in Cloud Connectors 14
- Cruise Control 15
- HBase 16
- HDFS 16
- Hive 22
- Cloudera Data Explorer (Hue) 24
- Cloudera Semantic Search 26
- What's New in Iceberg REST Catalog 27
- Iceberg 27
- Impala 30
- Kafka 36
- Knox 39
- Kudu 40
- Livy 41
- Navigator Encrypt 42
- Oozie 43
- Ozone 43
- Phoenix 53
- Ranger 54
- What's New in Ranger KMS 59
- Schema Registry 60
- Solr 60
- Spark 61
- Spark Atlas Connector 63
- Sqoop 63
- Streams Messaging Manager 63
- Streams Replication Manager 64
- YARN and YARN Queue Manager 65
- Zookeeper 67

Fixed Issues In Cloudera Runtime 7.3.2 68
- Atlas 68
- Avro 75
- Cloud Connectors 76
- Cruise Control 77

Fixed issues in Iceberg REST Catalog 77
HDFS 77
HBase 79
Hive 79
Cloudera Data Explorer (Hue) 90
Impala 94
Fixed Issues in Apache Iceberg 103
Kafka 111
Kudu 112
Knox 114
Livy 116
Navigator Encrypt 117
Oozie 117
Ozone 119
Parquet 138
Phoenix 139
Ranger 139
Ranger KMS 146
Schema Registry 147
Solr 148
Spark 148
Spark Atlas Connector 149
Sqoop 150
Streams Messaging Manager 150
Streams Replication Manager 152
Yarn and Yarn Queue Manager 153
ZooKeeper 158

Known Issues In Cloudera Runtime 7.3.2 161
Atlas 161
Avro 173
Known Issues in Cloud Connectors 173
Cruise Control 174
HBase 174
HDFS 175
Hive 176
Cloudera Data Explorer (Hue) 180
Iceberg 183
Known Issues in Iceberg REST Catalog 186
Impala 186
Kafka 192
Knox 197
Kudu 200
Livy 200
Navigator Encrypt 201
Oozie 201
Ozone 202
Parquet 209
Phoenix 210
Ranger 210
Ranger KMS 214
Schema Registry 215
Solr 216
Spark 217
Spark Atlas Connector 218

Sqoop....219
Streams Messaging Manager....220
Streams Replication Manager....222
Yarn and Yarn Queue Manager....223
ZooKeeper....226

Behavioral Changes in Cloudera Runtime 7.3.2....226
Behavioral Changes in Atlas....226
Behavioral Changes in Avro....228
Behavioral Changes in Cloud Connectors....228
Behavioral Changes in Cruise Control....229
Behavioral Changes in HBase....230
Behavioral Changes in HDFS....232
Behavioral Changes in Hive....233
Behavioral Changes in Cloudera Data Explorer (Hue)....237
Behavioral Changes in Impala....237
Behavioral Changes in Apache Iceberg....240
Behavioral Changes in Iceberg REST Catalog....240
Behavioral Changes in Kafka....240
Behavioral Changes in Apache Knox....242
Behavioral Changes in Kudu....243
Behavioral Changes in Livy....243
Behavioral Changes in Oozie....244
Behavioral Changes in Ozone....245
Behavioral Changes in Parquet....250
Behavioral Changes in Phoenix....250
Behavioral Changes in Apache Ranger....251
Behavioral Changes in Ranger KMS....254
Schema Registry....255
Behavioral Changes in Solr....256
Behavioral Changes in Spark....257
Behavioral Changes in Spark Atlas Connector....259
Behavioral Changes in Sqoop....259
Behavioral Changes in Streams Messaging Manager....260
Behavioral Changes in Streams Replication Manager....261
Behavioral Changes in YARN....263
Behavioral Changes in ZooKeeper....264

Release Matrix....266

Cloudera Runtime Component Versions....266
Cloudera Runtime 7.3.2.20000....266
Cloudera Runtime 7.3.2.10000....267
Cloudera Runtime 7.3.2.100....268
Cloudera Runtime 7.3.2....270

Using the Cloudera Runtime Maven repository....271
Cloudera Runtime 7.3.2.20000-258....271
Cloudera Runtime 7.3.2.10000-317....291
Cloudera Runtime 7.3.2.100-296....310
Cloudera Runtime 7.3.2.0-957....329

Fixed Common Vulnerabilities and Exposures..... 348
Cloudera Runtime 7.3.2.20000 SP2..... 348
Cloudera Runtime 7.3.2.10000 SP1..... 348
Cloudera Runtime 7.3.2.100 CHF 1..... 362
Cloudera Runtime 7.3.2..... 366

Deprecation Notices In Cloudera Runtime 7.3.2..... 376
Platform and OS..... 377
Deprecation Notices for Cruise Control..... 378
Deprecation Notices for Apache Impala..... 378
Deprecation Notices for Apache Kafka..... 378
Deprecation Notices for Apache Livy..... 379
Deprecation Notices for Apache Ozone..... 379
Deprecation Notices for Apache Ranger..... 379
Deprecation Notices for Apache Spark..... 380
Deprecation Notices for Streams Replication Manager..... 380
Deprecation Notices for Apache Zeppelin..... 380

API Modifications and Removals..... 381
Components with API differences in Cloudera Runtime 7.3.2.10000 SP1..... 381
API Compatibility changes in 7.3.2 for HBase..... 381
API Compatibility changes in 7.3.2 for ZooKeeper..... 381
Components with API differences in Cloudera Runtime 7.3.2.100 CHF 1..... 383
API Compatibility changes in 7.3.2.100 for ZooKeeper..... 383
Components with API differences in Cloudera Runtime 7.3.2..... 384
API Compatibility changes in 7.3.2 for Hadoop..... 384
API Compatibility changes in 7.3.2 for HBase..... 385
API Compatibility changes in 7.3.2 for Kafka..... 399
API Compatibility changes in 7.3.2 for Spark..... 400
API Compatibility changes in 7.3.2 for ZooKeeper..... 400

Unsupported Features..... 401
HDFS..... 401

Overview

This document provides a summary of the latest updates in Cloudera Runtime 7.3.2 and its Service Packs and Cumulative Hotfixes. It includes new features, improvements, known and fixed issues, technical previews, and more. For detailed, component-level information, see the Cloudera documentation.

Understanding Cloudera platform versioning

Starting with Cloudera 7.3.2.0, versioning follows the <Base>.<Major>.<Minor>.<Patch> format:

- **Base**: Indicates the platform base version.
- **Major**: Represents the major version.
- **Minor**: Represents the minor version.
- **Patch**: Identifies a Cumulative Hotfix (CHF) or Service Pack (SP). A value of 0 in this position denotes a standard minor release.

For example, in version **7.3.2.0**, **7** is the base version, **3** is the major version, **2** is the minor version, and **0** indicates that it is an initial minor release (not a CHF or SP). For more details, see the Cloudera Platform Support Policy.

Business value and strategic benefits of Cloudera 7.3.2.0

Cloudera 7.3.2.0 delivers three primary business benefits focused on cost efficiency, risk mitigation, and expanded operational capabilities:

Significant cost savings and infrastructure efficiency

Cloudera 7.3.2.0 release introduces powerful automation and optimization tools that maximize existing resources and lower operational expenditures:

- **Storage cost reductions:** The new Cloudera Storage Optimizer automatically identifies infrequently accessed data in Ozone and converts it to space-efficient erasure coding. This process reduces the storage footprint by 45% to 60%, effectively providing additional usable capacity without increasing licensed capacity.
- **Optimized compute costs:** Hive now supports ARM architecture (such as AWS Graviton), enabling you to run heavy workloads on significantly more cost-effective and energy-efficient hardware.
- **Automated table maintenance:** The Cloudera Lakehouse Optimizer provides automated table maintenance for Iceberg tables, which simplifies table management, accelerates query performance, and directly reduces operational overhead and costs.

Comprehensive risk reduction and regulatory compliance

Cloudera 7.3.2.0 release improves centralized authorization and cryptography to address business risks associated with data governance and security:

- **Unified cloud data governance:** Apache Ranger Remote Authorization Service (RAZ) now natively supports Amazon S3-compatible object stores. This enables administrators to apply consistent access policies to secure sensitive data across diverse cloud environments.
- **Hardened security baselines:** Cloudera has upgraded core components to JDK 17, utilizing stronger cryptographic algorithms to reduce security vulnerabilities and ensuring long-term patch support.
- **Enhanced FIPS compliance:** Ozone now supports fully FIPS-compliant SASL mechanisms (DIGEST-SHA), replacing legacy protocols (such as MD5 and DES) with SHA256 and AES to meet strict regulatory standards.

Operational scalability and future-proofing

The update delivers critical networking and high-availability features that unlock new use cases and ensure the platform scales with modern business demands:

- **Network scalability:** IPv6 dual-stack functionality now spans major components, including HBase, Hive, Impala, Kafka, Kudu, Phoenix, Cloudera Data Explorer (Hue), and ZooKeeper. This ensures seamless communication over both IPv4 and IPv6 to future-proof deployments.
- **Advanced disaster recovery:** Organizations can now leverage Apache Flink to replicate Kudu tables across clusters continuously, capturing only changed rows to support critical disaster recovery scenarios efficiently.
- **Simplified data modeling:** Support for one-dimensional arrays in Kudu and Impala allows data teams to store multiple related values in a single column. This improves query performance and unlocks new analytics use cases by eliminating resource-heavy joins.

What's new in Platform Support

Cloudera Runtime 7.3.2 release, its service packs (such as SP1), and its cumulative hotfixes introduce expanded platform support across Cloudera Manager 7.13.2, its service packs, and cumulative hotfix releases. Review these additions to plan your infrastructure upgrades and take advantage of newly supported operating systems, databases, Python versions, and Java Development Kits (JDKs).

Cloudera Runtime 7.3.2 SP2 (7.3.2.20000)

There are no new platform support enhancements in this release.

Cloudera Runtime 7.3.2 SP1 (7.3.2.10000)

Platform support enhancements

New OS support: Cloudera Runtime 7.3.2 SP1 release adds support for the following operating systems:

- RHEL 9.6 in FIPS mode
- SLES 15 SP7

For more information about the operating system support, see Cloudera Support Matrix.

New database support

The Cloudera Runtime 7.3.2 SP1 release adds support for the following database versions:

- PostgreSQL 18
- Oracle 26ai (also known as Oracle 23c or 23ai)

New Python support

The Cloudera Runtime 7.3.2 SP1 release adds support for the following Python version:

- Python 3.12

New JDK support

None

Cloudera Runtime 7.3.2 CHF 1 (7.3.2.100)

There are no new platform support enhancements in this release.

Cloudera Runtime 7.3.2

Platform support enhancements

New OS support: Cloudera Runtime 7.3.2 release adds support for the following operating systems:

- • RHEL 9.6
  - Rocky Linux 9.6

- SLES 15 SP6
- Ubuntu 24.04

For more information about the operating system support, see Cloudera Support Matrix.

New database support

Cloudera Runtime 7.3.2 release adds support for the following database versions:

- MariaDB 11.4

New JDK support

None

Runtime API Compatibility

Runtime API compatibility defines the architectural standards for communication between software components ensuring backward compatibility and consistent client-server interoperability as systems evolve.

Introduction to the API compatibility

API compatibility ensures seamless interaction between client and server versions by maintaining backward compatibility, and REST protocol consistency to prevent breaking changes across Cloudera Runtime releases.

The API defines a set of rules and definitions for the client-server communication. The client requests services or resources from the server, which processes the request and returns a response. API compatibility allows different client and server versions to interact successfully. Typically, a client supports its corresponding API version and lower backward-compatible versions.

Ensuring API compatibility has the following benefits:

- Changes are not broken: Avoid removing, renaming, or changing mandatory parameters within the software code.
- Versioning: Implement API versioning (for example, /version1/, /version2/) to introduce changes without affecting existing users and their usage patterns.
- Deprecation policy: Call out old APIs as deprecated before removing them.
- Validation opportunities: Verify changes that introducing new code does not break compatibility with older client versions.

Supported touchpoints for API compatibility

Seamless API integration enables different software systems to interact and exchange data automatically, to ensure real-time synchronization, reduced manual effort, and improved efficiency.

Cloudera supports the following touchpoints to ensure forward and backward- API compatibility across different Cloudera runtime releases:

- Rest APIs
- Command-Line Interfaces (CLIs)
- Client RPM / Debian Packages
- Client-Server Interaction (Atlas / Ranger / Data Hubs / Data Services)

Client-server compatibility

Client-server API compatibility in Cloudera Enterprise Data Platform ensures seamless interactions between components, even with updates to the client or the server.

Cloudera Enterprise Data Platform involves multiple layers for which API compatibility is critical. Each component in the Cloudera stack follows its own way of exchanging data using the API compatibility touchpoints, and the coverage is only being enhanced for the compatibility options.

Multiple Cloudera Runtime releases complement each other to maintain consistent API interactions.

The following table lists the supported Cloudera Runtime releases that are forward- and backward-compatible:

Table 1: Client-Server compatibility

Client	Server
7.3.2.0	7.3.2.0
7.3.1.500	7.3.2.0
CDH: 7.3.1.500-176	7.3.2.0
7.1.9.1063 and 7.1.9.1067 (Atlas, Hive, and Impala)	7.3.2.0
7.2.18.500	7.3.2.0
CDH: 7.2.18.500-67	7.3.2.0
7.2.17.1200	7.3.2.0
7.2.18.900-29	7.3.2.0
7.1.7.3023 (Hive, Ranger, Atlas, and Impala)	7.3.2.0
7.3.2.0	7.3.1.500/400
7.3.2.0	7.2.18.500, 7.2.18.900, and 7.2.18.1110
7.3.2.0	7.2.17.1200
7.3.2.0	7.1.9.1063 and 7.1.9.1067 (Atlas, Hive, and Impala)
7.3.2.0	7.1.7.3023 (Hive, Ranger, Atlas, and Impala)

For more information about each of these releases, refer to the Cloudera Product documentation.

Supported Cloudera Runtime components

Runtime components that are compatible with the categorised touchpoints.

The following table lists the Runtime components that are associated with the respective touchpoints for which the APIs are forward / backwards compatible between Cloudera Runtime releases.

Important:

- HDFS ACLs are a core filesystem feature implemented inside the NameNode and permission model, affecting how access checks are performed at runtime. While ACL operations are exposed through multiple interfaces (Java API, CLI, WebHDFS), the compatibility concern is primarily about internal permission semantics and enforcement behavior.
- Ozone S3 acts as a full-fledged S3-compatible service layer that enables external clients (like AWS SDKs or S3 tools) to communicate with Ozone storage. Ozone's S3 Gateway enables integration with a wide range of cloud-native and analytics applications.
- Ozone HTTPFS provides a REST interface to perform filesystem operations over HTTP, similar to WebHDFS. Ozone HTTPFS is essentially an API gateway that exposes filesystem functionality, without introducing a new protocol or ecosystem. Ozone HTTPFS is categorized under the REST / Thrift APIs category.

Table 2: Supported Runtime compatibility

Runtime component	REST / Thrift APIs	Client Libraries and SDKs	Client-Server Interaction (Atlas / Ranger / Data Hubs / Data Services)	Command-Line Interface
Atlas	Yes	Yes	Yes	N/A
Ranger	Yes	Yes	Yes	N/A
Ranger KMS	Yes	Yes	Yes	N/A
RAZ	Yes	Yes	Yes	N/A
Knox	Yes	Yes	Yes	N/A
HBase	N/A	Yes	N/A	Yes
HDFS	Yes	Yes	N/A	Yes
Phoenix	Yes	N/A	N/A	N/A
Ozone	Yes	Yes	Yes	Yes
Hive	Yes	Yes	N/A	Yes
Impala	Yes	Yes	Yes	Yes
Livy	Yes	N/A	N/A	N/A
Spark	Yes	Yes	Yes	Yes

What's New in Cloudera Runtime 7.3.2

This version of Cloudera Runtime provides you with several new capabilities. Learn how the new features and improvements benefit you.

Note:

For information about Cloudera Base on premises 7.3.2 service packs upgrade paths, see Supported in-place upgrade paths

What's New in Apache Atlas

New features and functional updates for Atlas are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

Impala lineage uses the Impala operation type

Impala lineage events include an explicit operationType field. Apache Atlas reads that value when it builds Impala process and column lineage so the lineage graph reflects the operation without relying only on parsing the full queryText. For queries that insert into a table or overwrite a table, the process segment in lineage qualified names uses INSERT or INSERT_OVERWRITE instead of QUERY or QUERY_WITH_CLAUSE.

For more information, see Impala lineage and Impala entities created in Atlas.

Atlas UI enhancements: React dashboard, contextual search, Create menu, and glossary import

The React-based Apache Atlas user interface adds an operational dashboard, a context-aware global search bar, a unified + Create entry point, and a single **Create Glossary** workflow that toggles between manual creation and file import.

- The **Overview** dashboard surfaces totals for entities and classifications, entity status distribution (Active, Shell, Deleted), classification types in use, entity type distribution, recently added entities, Kafka topic summary, and a **Recent Activity** stream for Created, Updated, Deleted, Purge, Import, and Export events.
- The global header includes primary + Create actions for Entities, Classifications, Glossaries, Business Metadata, and Enums.
- Quick search supports suggestions and a scope drop-down for Entities, Classifications, Terms (glossary), and Business Metadata.
- In **Create Glossary**, you can download import templates and import glossary terms by dragging and dropping .csv, .xls, or .xlsx files.

You can still switch between Classic and New UI experiences. The New UI search panel lists entity types, classifications, and glossary terms with filtering for empty service types, unused classifications, and category or term glossary views, plus flat tree navigation for entities and classifications. For more information, see Apache Atlas dashboard tour, Using Basic Search, and Creating glossaries.

Distributed notification processing framework

Apache Atlas introduces a two-phase distributed processing framework to handle high-throughput metadata ingestion from sources like Spark and Hive. This moves Atlas from a single-threaded sequential ingestion pipeline to a parallel pipeline using Kafka topics. The preprocessor routes events to ATLAS_METADATA_* and ATLAS_LINEAGE_* topics, which are then processed concurrently by multiple consumers. This significantly improves horizontal scalability while preserving referential integrity by deterministically routing related entities to the same topic. For more information, see Configuring distributed notification processing.

TLS 1.3 support

Apache Atlas supports TLS 1.3 for stronger security and compliance. In Cloudera Manager, advanced TLS protocol and cipher options, including inheritance from global TLS settings, are available in the Atlas service configuration. Cloudera recommends upgrading to TLS 1.3 when possible. For more information, see Configuring TLS/SSL for Apache Atlas.

Header-based authentication

Apache Atlas supports header-based authentication so a trusted gateway can pass identity using HTTP headers. Default values align with x-awc-username, x-awc-roles, and x-awc-requestid. Configure atlas.authn.header.* keys in atlas-application.properties. For more information, see Configuring header-based authentication for Apache Atlas.

Optional PostgreSQL backend for JanusGraph and entity audit storage

Atlas can optionally use PostgreSQL as the JanusGraph KeyColumnValue persistence backend (atlas.graph.storage.backend=rdbms and related JDBC settings) to reduce footprint when you do not require a horizontally scalable HBase-backed graph store. You can also store entity audit events in a relational database by setting atlas.EntityAuditRepository.impl to org.apache.atlas.repository.audit.rdbms.RdbmsBasedAuditRepository. Default installations continue to use the existing HBase unless you select PostgreSQL during setup.

Cloudera Manager now provides a native Backend Selection step during the Atlas **Add Service** wizard to simplify configuring the PostgreSQL connection, including support for SSL/TLS-encrypted JDBC connections. For more information, see Configuring a PostgreSQL database for Atlas and Installing Atlas using Add Service.

Faster Atlas entity detail for large relationship graphs

The first load of an entity detail page omits large embedded relationship payloads so the page opens faster. On the Relationships tab, you can switch between Table mode (a card per relationship type with counts, sort, optional deleted relationships, and paged loads as you scroll) and Graph mode. The relationship search API supports pagination, optional sorting, and more reliable counts for very large graphs. Together, these changes address cases such as Apache Ozone buckets or wide Hive tables for which the detail view previously failed to display Properties, Classifications, Relationships, or Audits because the response was too large.

For more information, see Apache Atlas dashboard tour.

Cloudera Runtime 7.3.2.100 CHF 1

Faster Atlas entity detail for large relationship graphs

The first load of an entity detail page omits large embedded relationship payloads so the page opens faster. On the Relationships tab, you can switch between Table mode (a card per relationship type with counts, sort, optional deleted relationships, and paged loads as you scroll) and Graph mode. The relationship search API supports pagination, optional sorting, and more reliable counts for very large graphs. Together, these changes address cases such as Apache Ozone buckets or wide Hive tables where the detail view previously failed to show Properties, Classifications, Relationships, or Audits because the response was too large.

For more information, see Apache Atlas dashboard tour.

Impala lineage uses the Impala operation type

Impala lineage events include an explicit operationType field. Apache Atlas reads that value when it builds Impala process and column lineage so the lineage graph reflects the operation without relying only on parsing the full queryText. For queries that insert into a table or overwrite a table, the process segment in lineage qualified names uses INSERT or INSERT_OVERWRITE instead of QUERY or QUERY_WITH_CLAUSE.

For more information, see Impala lineage and Impala entities created in Atlas.

Cloudera Runtime 7.3.2:

Cloudera Runtime 7.3.2 introduces new features of Atlas and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see New Features.

New React-based user interface for Apache Atlas

Apache Atlas now features a redesigned React-based user interface (UI) that offers enhanced usability and streamlined metadata management. You can switch between the Classic and New UI experiences. The new interface introduces an improved search panel that automatically lists all available entity types, classifications, and glossary terms, with one-click access to relevant members. Enhanced filtering capabilities allow you to display empty service types, unused classifications, and toggle between category or term views in the glossary. Additionally, the interface displays entities and classifications in a collapsed flat tree view to simplify navigation of complex metadata hierarchies.

For more information, see Apache Atlas dashboard tour.

Apache Atlas component upgraded to 2.4.0

The Atlas runtime component is upgraded from 2.1.0 to 2.4.0. Several stability and correctness fixes are included from the upstream release for bugs, including UI improvements for classification propagation settings.

Atlas auto-purging introduced



Important: This feature is available as technical preview and is under entitlement. To obtain the required entitlement, contact your Cloudera Account Representative.

The automated entity auto-purging feature addresses potential performance and storage issues caused by the manual purge strategy. Previously, Atlas preserved metadata by only marking entities as deleted (soft-delete). This led to query performance degradation and increased storage usage as soft-deleted entities accumulated. While these entities could be manually deleted by using the `PUT /admin/purge/ API` call, the process failed to remove column lineage entities associated with soft-deleted process entities.

The new, cron-based system can clean up obsolete process entities, including their column lineage entities, that are no longer relevant. This prevents sparse graphs and significantly improves metadata hygiene and query performance. For more information, see Atlas Auto-Purging overview.

Support replication of Atlas data from on-premises to on-premises



Important: This feature is available as technical preview and is under entitlement. To obtain the required entitlement, contact your Cloudera Account Representative.

Atlas now supports asynchronous import of metadata using Kafka. Previously, the system relied on synchronous imports, which required the HTTP connection to remain open until the entire process completed. This caused timeouts for large datasets and making concurrent imports fragile. With asynchronous import, a client submits an import request that is immediately staged and queued as a Kafka message, and receives an import ID in response without waiting for processing to finish. Atlas processes the import in the background and persists the request state, including received time, processing start time, completion time, and outcome.

The following new REST API endpoints are available:

- POST /api/atlas/admin/async/import — submit an asynchronous import; returns immediately with an import ID
- GET /api/atlas/admin/async/import/status — list all async import statuses
- GET /api/atlas/admin/async/import/status/{importId} — get the status of a specific import
- DELETE /api/atlas/admin/async/import/{importId} — abort a specific queued import

Atlas upgraded to use JDK 17

Atlas now runs on Java 17, upgraded from Java 8. JDK 17 is a Long-Term Support (LTS) release that brings improved performance, enhanced security, and better long-term maintainability. JDK 17 has the following key benefits for Atlas users:

- Improved garbage collection, resulting in lower latency and more efficient memory usage for metadata-intensive workloads.
- Stronger cryptographic algorithms reducing security vulnerabilities.
- Long-term support guaranteed until at least 2029, ensuring continued security patches.

Logback introduced as logging framework

Apache Atlas now uses Logback as its logging framework, replacing Log4j2. This change enhances security and simplifies log management. It also enables you to add any new properties overriding existing properties.

The new logging framework provides simplified configuration by using native XML configuration instead of the .properties file.

- Go to Cloudera Manager Atlas Server XML Override to replace the complete configuration file.

Configuration values for file size and rotation remain the same.

What's New in Cloud Connectors

Learn about the new features of Cloud Connectors in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

Rebase to Hadoop 3.4.2

This release of the Cloud Connectors is based on Hadoop 3.4.2. See the following upstream resources for more information on the changes:

- Apache Hadoop 3.4.2
- Apache Hadoop 3.4.1
- Apache Hadoop 3.4.0

What's New in Cruise Control

New features and functional updates for Cruise Control are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

Centralized TLS protocol and cipher settings

When you enable the TLS_ADVANCED_CONTROL feature flag, the webserver.ssl.include.protocols and webserver.ssl.include.ciphers Cruise Control webserver properties default to the centralized Cloudera Manager TLS settings. Do not configure these properties in the safety valve when the flag is enabled. For information about enabling the feature flag, see *Transport Layer Security Support and Custom Protocol Configuration*.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

New configuration parameter for controlling IP stack preference

A new cc.additional.java.options configuration parameter is available on the Cruise Control configuration page in Cloudera Manager. The default value sets the IP protocol to IPv4.

New intra.broker.goals configuration for Cruise Control

Cloudera Manager introduces a new intra.broker.goals configuration for Cruise Control. The default value includes com.linkedin.kafka.cruisecontrol.analyzer.goals.IntraBrokerDiskCapacityGoal and com.linkedin.kafka.cruisecontrol.analyzer.goals.IntraBrokerDiskUsageDistributionGoal.

This has an effect on the existing Default Goals (default.goals) configuration, which must be a subset of Supported Goals and Supported Intra Broker Goals.

Additionally, the intra.broker.goals configuration no longer needs to be defined in an advanced configuration snippet if done previously.

What's New in Apache HBase

New features and functional updates for HBase are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

TLS 1.3 support for HBase

HBase supports Transport Layer Security (TLS) 1.3, providing stronger encryption, faster handshakes, and improved security over TLS 1.2. TLS 1.3 is natively available across all HBase communication channels, including the HBase Master UI, RegionServer UI, REST Server, and Thrift Server.

For more information, see Configuring the TLS versions for HBase

Cloudera Runtime 7.3.2.100 CHF1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of HBase and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see New Features.

IPv6 support for HBase

Starting with the 7.3.2 release, HBase client supports IPv6 with dual-stack functionality, allowing seamless communication over both IPv4 and IPv6 networks. This capability improves network scalability, future-proofs deployments, and enhances overall platform security.

For more information, see Enabling IPv6 support for HBase.

HBase JDK 17 upgrade

HBase only supports JDK 17 starting with the 7.3.2 release. Upgrade to JDK 17 to gain improved performance, increased security, modern language features, and long-term support, ensuring your application remains competitive and maintainable.

Apache HBase component upgraded to 2.6.3

The HBase runtime component is upgraded from 2.4.17 to 2.6.3. This release incorporates stability and correctness fixes from the upstream, which deliver performance and maintenance improvements, and enhanced client connectivity support.

What's New in HDFS

Learn about the new features of HDFS in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new features in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features in this release.

Cloudera Runtime 7.3.2

Support for G1 Garbage Collector (G1GC) with JDK 17

When running HDFS on JDK 17, all HDFS server processes now automatically use the G1 Garbage Collector (G1GC). No action is required for most deployments.

Hadoop rebase summary

In Cloudera Runtime 7.3.2, Apache Hadoop is rebased to version 3.4.1. The Apache Hadoop upgrade improves overall performance and includes all the new features, improvements, and bug fixes from versions 3.2, 3.3, and 3.4.

Table 3: New features added from Apache Hadoop 3.2 to 3.4 versions

Apache Hadoop version	Apache Jira	Name	Description
3.2	HDFS-10285	Storage Policy Satisfier in HDFS	StoragePolicySatisfier (SPS) allows you to track and fulfill the storage policy requirement of a given file or directory in HDFS. You can specify a file or directory path for SPS to evaluate by running the `hdfs storagepolicies -satisfyStoragePolicy -path <path>` command or invoking the `HdfsAdmin#satisfyStoragePolicy(path)` API. If blocks have storage policy mismatches, SPS moves the replicas to a different storage type to fulfill the storage policy requirements. Because API calls go to the NameNode to track the invoked satisfier path (iNodes), you must enable SPS on the NameNode by setting the dfs.storage.policy.satisfier.mode property to exte rnal in the hdfs-site.xml file. You can be disable the configuration dynamically without restarting the NameNode. SPS must be started outside the NameNode using the hdfs –da emon start sps command. To run the Mover tool explicitly, you must disable the SPS first. For more information, see the *Storage Policy Satisfier (SPS)* section in the Archival Storage . By default, Storage Policy Satisfier is disabled in Cloudera HDFS.
3.3	HDFS-14118	DNS resolution for nameservices to IP addresses	HDFS clients can use a single domain name to discover servers such as namenodes, routers, or observers instead of explicitly listing out all hosts in the configurations.
3.3	HDFS-13783	Long-running service process for the Balancer	Adds the -asService parameter to the Balancer CLI. This parameter enables the Balancer to run as a long-running service process for easier monitoring. By default, this service mode is disabled in Cloudera HDFS.
3.3	HADOOP-16398	Hadoop metrics export to Prometheus	If the hadoop.prometheus.endpoint.enabled configuration is set to true, Prometheus-friendly formatted metrics can be obtained from the /prom endpoint of Hadoop daemons. The default value is false. By default, the configuration is disabled in Cloudera HDFS.

Apache Hadoop version	Apache Jira	Name	Description
3.3	HDFS-13571	Dead node detection	When an unresponsive node blocks a DFSInputStream, the system detects the failure and shares this information to other streams in the same DFSClient. This prevents the remaining streams from attempting to read from and be blocked by the dead node .
3.4	HADOOP-17010	Queue capacity weights in FairCall Queue	When the FairCallQueue feature is enabled, you can specify capacity allocation among all sub-queues using the ipc.<port>.callqueue.capacity.weights configuration. The value of this configuration is a comma-separated list of positive integers, each of which specifies the weight associated with the sub-queue at that index. The number of weights in this list must match the IPC scheduler priority levels defined in scheduler.priority.levels. By default, each sub-queue is associated with a weight of 1, meaning all sub-queues are allocated with the same capacity. This configuration is optional in Cloudera HDFS.
3.4	HDFS-13183	Standby NameNode processing for the getBlocks requests to reduce the active load	Enables the balancer to redirect getBlocks request to a standby NameNode, thus reducing the performance impact of the balancer to the active NameNode. The feature is disabled by default. To enable it, set the dfs.ha.allow.stale.reads configuration of the balancer to true in the hdfs-site.xml file.
3.4	HDFS-15025	NVDIMM storage media support for HDFS	Adds the NVDIMM storage type and the ALL_NVDIMM storage policy for HDFS. The NVDIMM storage type is for non-volatile random-access memory storage media whose data survives when the DataNode restarts.
3.4	HDFS-15098	SM4 encryption method for HDFS	The SM4/CTR/NoPadding encryption codec is now added. It requires OpenSSL 1.1.1. or higher version for native implementation.
3.4	HDFS-15747	RBF: Renaming across sub-namespaces	Renames multiple namespaces on the Federation balance tool.
3.4	HDFS-15547	Dynamic disk-level tiering	Archival Storage allows the configuration of DISK and ARCHIVE types on the same device. It enables optimizing disk I/O in heterogeneous environments.
3.4	HDFS-16595	Slow peer metrics	NameNode metrics that represent slownode JSON now include three important factors, namely median, median absolute deviation, and upper latency limit, which can help you determine how urgently a given slownode requires manual intervention.

Table 4: Improvements added from Apache Hadoop 3.2 to 3.4 versions

Apache Hadoop version	Apache Jira	Name	Description
3.3	HDFS-7133	Clearance of the namespace quota on /	Namespace quota on root / can be cleared now.
3.3	HADOOP-13363	Upgrade of protobuf from 2.5.0 to the highest version	Upgrade protobuf from 2.5.0 to the highest version.
3.3	HADOOP-16670	Removal of Submarine code from Hadoop codebase	The Submarine subproject has been moved to its own Apache Top Level Project. Therefore, the Submarine code has been removed from Hadoop 3.3.0 and higher versions. For more information, see https://submarine.apache.org/.

Apache Hadoop version	Apache Jira	Name	Description
3.3	HADOOP-17021	Addition of the -concat fs command	The `hadoop fs` utility includes a `-concat` command available on all filesystems that support the `concat` API, including HDFS and WebHDFS.
3.4	HADOOP-17222	Creation of socket address leveraging the URI cache	The DFS client can use the newly added URI cache when creating a socket address for read operations. When enabled, creating a socket address will use cached URI object based on host:port to reduce the frequency of URI object creation. To enable it, set the following configuration key to true: `<property>   <name>dfs.client.read.uri.cache.enabled</name>   <value>true</value> </property>` By default, the configuration is disabled in Cloudera HDFS.
3.4	HDFS-15814	Configurable parameters for DataNodeDiskMetrics	Makes certain parameters configurable for DataNode DiskMetrics.
3.4	HADOOP-16747	Support for Python 3	Supports Python 3.
3.4	HADOOP-17514	Removal of the `trace` subcommand from the Hadoop CLI	The `trace` subcommand of the Hadoop CLI has been removed as a follow-up of removal of the TraceAdmin protocol.
3.4	HADOOP-16524	Automatic keystore reloading for HttpServer2	Adds auto-reload of the keystore. Adds the following new configuration (default is 10 seconds): `ssl.{0}.stores.reload.interval` The refresh interval is used to check if either the truststore or keystore certificate file has changed.
3.4	HDFS-15975	Replacement of LongAdder with AtomicLong	This update changes public fields in DFSHedgedReadMetrics. If you are using the public member variables of DFSHedgedReadMetrics, you must access them using the public API.
3.4	HADOOP-17956	Replacement of all default Charset usage with UTF-8	All of the default charset usages have been replaced to UTF-8. If the default charset of your environment is not UTF-8, the behavior can be different.
3.4	HDFS-16278, HDFS-16285, and HDFS-16419	Cross-platform support for HDFS snapshot tools	Supports HDFS snapshot tools cross-platform.
3.4	HDFS-15382	Division of the single FsDatasetImpl lock into volume-grained locks	Throughput is one of the core performance evaluations for the DataNode instance. However, it does not reach the best performance, especially for Federation deployments, all the time because of the global coarse-grain lock. HDFS-16534, HDFS-16511, HDFS-15382 and HDFS-16429 aim to split the global coarse-grain lock into a fine-grain lock, which is a double-level lock for blockpool and volume, to improve the throughput and avoid lock impacts between blockpools and volumes.

Apache Hadoop version	Apache Jira	Name	Description
3.4	HADOOP-18188	Support for the `touch` command for directories	Previously, `hadoop fs -touch` command threw a PathIsDirectoryException error for a directory. Now, the command supports directories and will not throw that exception.
3.4	HADOOP-13332	Removal of Jackson 1.9.13 and migration to the 2.x codebase	Remove Jackson 1.9.13 and switch all Jackson code to the 2.x codebase.
3.4	HADOOP-18332	Removal of the rs-api dependency by downgrading Jackson to 2.12.7	Downgrades Jackson from 2.13.2 to 2.12.7 to fix class conflicts in downstream projects. This version of Jackson does contain the fix for CVE-2020-36518.
3.4	HADOOP-18442	Removal of the hadoop-openstack module	The swift:// connector for Openstack support has been removed because of fundamental limitations. such as Swift's handling of files greater than 4 GB. Because most object store services now export a subset of the S3 protocol, use the s3a connector instead. The hadoop-openstack JAR remains as an empty artifact to prevent build failures for projects that still declare the JAR as a dependency.
3.4	HADOOP-16206	Migration from commons-logging to reload4j	Migrates commons-logging to reload4j.
3.4	HADOOP-17524	Removal of EventCounter and Log counters from JVMMetrics	Removes the log counter from JVMMetrics because this feature strongly depends on the Log4J 1.x API.
3.4	HADOOP-18206	Cleanup of commons-logging references in the codebase	Cleans up the commons-logging references in the codebase.
3.4	HADOOP-18820	AWS SDK v2: Optional v1 bridging support	The v1 aws-sdk-bundle JAR is removed. It is only required for third party applications or when using v1 SDK AWSCredentialsProvider classes. Standard providers automatically migrate from the v1 to v2 classes. Therefore this change only affects third-party providers or environments using very esoteric classes in the V1 SDK . For more information, see Upgrading S3A to AWS SDK V2.
3.4	HADOOP-18876	ABFS: Default buffer change from disk to bytebuffer for fs.azure.data.blocks.buffer	The default value for the fs.azure.data.blocks.buffer configuration is now changed from disk to bytebuffer. This speeds up writing to Azure storage, at the risk of running out of memory, especially if many threads are writing to ABFS at the same time and the upload bandwidth is limited. If jobs do run out of memory while writing to ABFS, change the option back to disk.
3.4	HADOOP-18993	S3A: Addition of the fs.s3a.classloader.isolation option	To load custom implementations of AWS Credential Providers through user-provided jars, set the {{fs.s3a.extensions.isolated.classloader}} configuration to {{false}}.

Apache Hadoop version	Apache Jira	Name	Description
3.4	HADOOP-19120	ABFS: Adaptation of ApacheHttpClient as a network library	Apache HttpClient 4.5.x is a new implementation of HTTP connections, supporting a large, configurable pool of connections along with the ability to limit their lifespan. Configure the networking library using the fs.azure.networking.library configuration option. The following are the supported values: • JDK_HTTP_URL_CONNECTION Uses the JDK networking library [Default] • APACHE_HTTP_CLIENT Uses Apache HttpClient **Important:** If you switch the networking library to the Apache HttpClient, you must ensure that Apache HttpCore and HttpClient are on the classpath.

Table 5: Issues fixed between Apache Hadoop versions 3.2 to 3.4

Apache Hadoop version	Apache Jira	Name	Description
3.3	HDFS-14396	Image loading failure from FSImageFile during downgrades from 3.x to 2.x	During a rolling upgrade from Hadoop 2.x to 3.x, NameNode could not persist erasure coding information, and therefore a user could not start using the erasure coding feature until the finalize operation was complete.
3.3	HADOOP-15358	SFTPConnectionPool connection leakage	SFTPConnectionPool connection leakage occured.
3.3	HDFS-14845	Bypass of `AuthenticationFilterInitializer` for `HttpFSServerWebServer`	HttpFS previously ignored standard Hadoop authentication settings. HttpFS now honors hadoop.http.authentication.* configurations and the previous httpfs.authentication.* configurations are deprecated. If both configurations are set, the deprecated httpfs.authentication.* configurations take precedence for backward compatibility.
3.4	HDFS-15380	RBF: Fetch failure for the real remote IP in `RouterWebHdfsMethods`	RBF could not fetch the real remote IP in `RouterWebHdfsMethods`.
3.4	HDFS-15281	ZKFC host binding fallback from the dfs.namenode.rpc-bind-host to dfs.namenode.rpc-address	ZKFC previously ignored the dfs.namenode.rpc-bind-host configuration. ZKFC now correctly binds the host address using dfs.namenode.servicerpc-bind-host if configured, falling back to dfs.namenode.rpc-bind-host. If neither configuration exists, ZKFC binds to the NameNode RPC server address (acting as dfs.namenode.rpc-address configuration).
3.4	HADOOP-18237	Apache Xerces Java upgrade to 2.12.2	An older version of Apache Xerces Java contained vulnerabilities. Apache Xerces has been updated to 2.12.2 to resolve CVE-2022-23437.
3.4	HADOOP-18621	CryptoOutputStream::close leakage	A resource leak occurred in CryptoOutputStream::close when encountering encrypted zones combined with quota exceptions occur.

Apache Hadoop version	Apache Jira	Name	Description
3.4	HADOOP-18329	Missing support for Semeru OE JRE 11.0.15.0	Support was missing for Semeru Runtimes, in which, due to vendor name-based logic and changes in the java.vendor property, failures could occur on Java 11 Runtimes 11.0.15.0 and higher versions.
3.4	HADOOP-19221	S3A: Recovery failure for S3A multipart block upload attempts Status Code: 400; Error Code: RequestTimeout	S3A upload operations failed to recover when the store returned a 500 error. These operations now successfully recover. You can control this retry behavior for 50x errors, excluding 503 throttling events, which remain independently processed, using the fs.s3a.retry.http.5xx.errors property. Th default value is true.

Hadoop interface standards and API changes

Hadoop provides many compatibility promises, including the following guarantees:

- Java APIs marked Stable must not change within a major release.
- Hadoop wire protocols are classified as Private and Stable, and remain forward-compatible across minor releases within a specific major version.
- The exposed Hadoop REST APIs are generally classified as Public and Evolving. However, within a specific API version number, they are treated as Public and Stable, meaning no incompatible changes are permitted.

The following library changes affect Cloudera components:

- The commons-lang (2) has been migrated to commons-lang3, which uses a new package name. Dependent Cloudera components must update their imports.
- Guava and some instances of Protobuf now come as shaded JARs from hadoop-thirdparty.
- A new bouncycastle version, including Java 17 and 21 class variants, is now available. Tools such as Gradle and maven-shade-plugin must be updated to higher versions or excluded when shading.
- Switched from javax.activation to jakarta.activation.
- org.apache.kerby.kerb-simplekdc is no longer imported by Hadoop, so components that depend on it must explicitly include it.
- Migrated from org.codehaus.jackson (1.x) to com.fasterxml.jackson (2.x).
- The Hadoop `EventCounter` class was removed in HADOOP-17524 with no replacement because it relied on deprecated log4j 1.x functionalitynot present in 2.x.
- Removed commons-logging. Hadoop now exclusively configures and uses reload4j now.

The following Java API changes affect Cloudera components:

- HADOOP-16645 – Switches S3A APIs to return a StoreContext rather than provide direct access to S3AFileSystem.
- HADOOP-18206 – Removed the IOUtils#cleanup method from Public-Evolving API as part of the migration away from commons-logging.
- HDFS-16982 – Renamed MutableQuantiles.quantiles to MutableQuantiles.QUANTILES.

What's New in Apache Hive

New features and functional updates for Hive are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

Hive Metastore leader election

You can now configure Hive Metastore (HMS) to elect a single leader instance that runs housekeeping tasks such as compaction and partition management. Leader election uses Hive locks to elect the leader automatically and to fail over when the leader stops responding, which prevents multiple HMS instances from duplicating the same work and overloading the metastore database. For more information, see .

Apache Jira: HIVE-26509

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Hive and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see New Features.

Hive user experience enhancements

Cloudera now provides several Hive user experience enhancements:

- Improved error handling for the `SHOW PARTITIONS` command: The `SHOW PARTITIONS` command now returns a concise execution error instead of a full stack trace when you run it against a non-partitioned table. This change simplifies the output and helps you identify configuration issues more quickly. (HIVE-26926)
- Enhanced error messages for the STORED BY clause: The error message displayed when you provide an invalid identifier or literal in the STORED BY clause is now more informative. This improvement provides better guidance if you mistakenly use STORED BY instead of STORED AS or provide an incorrect storage format. (HIVE-27957)
- Enhanced task attempt log clarity: The task attempt log now includes explicit context for boolean values to improve readability. You can now clearly identify whether a task attempt is guaranteed or not within the log messages. (HIVE-28246)
- Enabled vectorized mode support for custom UDFs: You can now use custom User-Defined Functions (UDFs) in vectorized execution mode. This enhancement ensures that custom functions, such as TIME_PARSE, integrate correctly with vectorized query plans to improve processing performance. (HIVE-28830)
- Resolved NullPointerException in TezSessionPoolManager: A NullPointerException (NPE) that occurred in the TezSessionPoolManager when the resource plan was null is now resolved. This fix improves the stability of the Tez session pool when updating triggers from an active resource plan. (HIVE-29007)

Apache Jira: HIVE-26926, HIVE-27957, HIVE-28246, HIVE-28830, HIVE-29007

Dropping Hive Metastore statistics

Statistics associated with tables, partitions, and columns in the Hive Metastore (HMS). This feature is particularly useful during migration or replication processes where large volumes of statistical data—generated for every table, partition, and column combination can significantly increase copy times. Removing unnecessary statistics, you can improve operational efficiency and reduce data transfer overhead.

Apache Jira: HIVE-28655

Enhanced Hive Metastore notification fetching with table filters

The Hive Metastore (HMS) notification fetch API now supports optional database and table name filters. This enhancement allows external engines and clients to fetch pending events specifically for a given table or list of tables, rather than scanning all notifications. Additionally, an index on the table name is now included in the HMS notification log to optimize query performance and

prevent table scans. This change enables more efficient metadata synchronization and improves performance for queries involving multiple tables.

Apache Jira: HIVE-27499

New command to display HiveServer2 and Hive Metastore connections

You can now use the `SHOW PROCESSLIST` command to display active operations and connection details for HiveServer2 (HS2) and Hive Metastore (HMS) instances. This feature provides a view of current sessions, including user names, IP addresses, query IDs, and execution states, similar to the process list functionality in MySQL. This command helps you troubleshoot stuck queries, monitor service load, and identify inappropriate connections for termination.

Apache Jira: HIVE-27829

Upgrading Calcite

Hive has been upgraded to Calcite version 1.33. This upgrade introduces various query optimizations that can improve query performance.

Hive on ARM Architecture

Hive is now fully supported on ARM architecture instances, including AWS Graviton and Azure ARM. This enables you to run your Hive workloads on more cost-effective and energy-efficient hardware.

ZooKeeper SASL authentication for Hive clients

You can now configure Hive clients to authenticate with a ZooKeeper ensemble that enforces Simple Authentication and Security Layer (SASL). By using the new hive.zookeeper.client.sasl.enforce property, `JDBC` and HS2 clients can successfully establish sessions in Kerberized environments that require secure service discovery and locking. For more information, see

Impala now supports OAuth Authentication

Cloudera now provides support for OAuth authentication using OAuth JWT bearer tokens. For more information, see

What's New in Cloudera Data Explorer (Hue)

New features and functional updates for Cloudera Data Explorer (Hue) are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

Data Explorer supports TLS 1.3 for enhanced security and compliance.

In Cloudera Manager, advanced TLS protocol and cipher options, including inheritance from global TLS settings, are available in the Data Explorer service configuration. Cloudera recommends upgrading to TLS 1.3 for improved security and protocol support.

For more information, see Configuring TLS/SSL for Data Explorer .

Data Explorer support for Python 3.12 on Red Hat Enterprise Linux (RHEL) 10.1.

Data Explorer supports Python 3.11 and Python 3.12 on Red Hat Enterprise Linux (RHEL) 10.1 starting in release 7.3.2.10000. Upgrade your Python environment to Python 3.11 or Python 3.12 to maintain compatibility with Data Explorer services.

Data Explorer SQL AI Enhancement:

The SQL AI Assistant introduces the following enhancements:

- Personally Identifiable Information (PII) filtering: Configures PII entity detection by using the pii_entities_to_detect configuration file to prevent sensitive data from being sent to the model.
- Automatic row limit caps: Automatically appends a default LIMIT 1000 clause to generate read queries in the SQL Editor user interface preview before running them to protect cluster performance.

For more information, see Cloudera Data Explorer (Hue) SQL AI user input validations.

Data Explorer SQL AI Assistant API (Technical Preview)

Previously, the SQL AI Assistant features in Data Explorer were accessible only through the Data Explorer user interface. You can now access the SQL AI Assistant through API. This enhancement enables programmatic SQL operations, including sending natural language requests directly to the API endpoints.

For more information, see Cloudera Data Explorer (Hue) SQL AI user input validations.

New Data Explorer Storage Browser (General Availability)

The Data Explorer Storage Browser is a web-based interface designed to provide seamless interaction with multiple file systems. With enhanced usability and functionality, the Storage Browser improves data management within the Hadoop ecosystem, offering a streamlined experience.

For more information, see Hue Storage Browser.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Cloudera Data Explorer (Hue) and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see New Features.

Product Branding Update

The product component previously known as **Hue** is now renamed to Cloudera Data Explorer (Hue). This change reflects UI and an updated branding initiative and will be rolled out in phases.

As part of this release, you can notice the following changes:

- A new logo in the UI
- The service name updated to Data Explorer in the UI
- The new product name reflected in documentation

Some UI references might still display the previous name as the branding update is completed incrementally in future releases.

No functional impact is associated with this change. All existing configurations, workflows, and integrations continue to work as before.

Data Explorer JDK 17 upgrade

Data Explorer supports JDK 17. Upgrading to JDK 17 ensures your application stays competitive and maintainable by providing improved performance, increased security, modern language features, and long-term support.

Data Explorer IPv6 support

Data Explorer supports IPv6 with dual-stack functionality, allowing seamless communication over both IPv4 and IPv6 networks. This capability improves network scalability and enhances overall platform security.

For more information, see Enabling Dual-Stack support for Data Explorer.

Python 3.11 support in Data Explorer

Data Explorer now supports Python 3.11 as the only supported version across all operating systems. Python 3.8, 3.9, and 3.10 are no longer supported.

Enhanced session security for Data Explorer

Data Explorer now includes security for the session ID (sessionid) cookie. This enhancement helps prevent unauthorized access, resulting in data exposure, unauthorized query execution, and job submission across connected Data Explorer services.

For more information, see Securing Data Explorer sessions.

Improved navigation pane interpreter visibility

Previously, enabling an application such as SparkSQL automatically displayed every associated default interpreter, which often cluttered the UI. Data Explorer now prioritizes a streamlined interface by no longer displaying all available interpreters by default. If your workflow requires the previous behavior, you can revert to displaying all available interpreters by enabling the enable_all_interpreters configuration flag.

For more information, see Enabling multiple editors in Cloudera Data Explorer (Hue).

Global Hive JDBC URL for Oozie workflows

Data Explorer now supports a global configuration for Hive JDBC connection strings used in Oozie workflow. Data Explorer derived these URLs from the default Hive configuration, requiring manual, per-workflow overrides for specialized setups such as High Availability (HA) quorums. You can now use the oozie_hs2_jdbc_url configuration property to set a single, cluster-wide JDBC URL.

For more information, see Configuring global Hive JDBC URL for Oozie.

What's New in Cloudera Semantic Search

New features and functional updates for Cloudera Semantic Search that are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new features in this release.

Cloudera Runtime 7.3.2.100 CHF1

Cloudera Semantic Search component upgraded to 3.6.0

The Cloudera Semantic Search component is upgraded from 2.18.0 to 3.6.0. This release incorporates upstream improvements that deliver enhanced performance, improved stability, and expanded functionality for search workloads.

JDK 21 bundled with Cloudera Semantic Search parcel

OpenSearch 3.x requires JDK 21, while Cloudera Runtime uses JDK 17. To address this requirement, JDK 21 is bundled directly inside the Cloudera Semantic Search parcel and is used exclusively by the OpenSearch process. This approach ensures compatibility with OpenSearch 3.x without affecting the Cloudera Runtime Java environment.

Document-Level Security for Cloudera Semantic Search

You can now restrict access to specific documents within an OpenSearch index based on user roles defined in Apache Ranger. Document-Level Security (DLS) enables administrators to enforce fine-

grained access control that goes beyond index-level permissions, ensuring that users can view only the documents they are authorized to access.

For more information, see Configuring Document-Level Security for Cloudera Semantic Search.

Cloudera Runtime 7.3.2

Cloudera Semantic Search [Technical Preview]

Cloudera Semantic Search is available as a technical preview. For more information about this service, see the *Cloudera Semantic Search service overview* documentation.



Important: This feature is not intended for production environments. To provide feedback, log a case through the Cloudera Support Portal. Cloudera does not guarantee troubleshooting or fixes for technical preview services.

What's New in Iceberg REST Catalog

New features and functional updates for Cloudera Iceberg REST Catalog are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2:

On-premises Data Sharing with Cloudera Iceberg REST Catalog on Ozone

Starting with Cloudera Runtime 7.3.2.20000, you can share Iceberg tables stored on Cloudera Object Store (powered by Apache Ozone) in Cloudera on premises environments. Setup uses the Hive Metastore embedded Cloudera Iceberg REST Catalog, Metering V2, Knox IDBroker, Ozone STS, and Knox topologies cdp-share-management (external client administration) and cdp-datashare-access (OAuth and Iceberg REST access for data consumers).

External clients register through Knox and Ranger. The REST Catalog vends temporary Ozone S3 credentials through the S3 Gateway so compliant engines can read shared table data. Iceberg tables whose storage locations use s3a:// paths are shareable; tables registered with ofs:// locations only cannot be shared with vended credentials.

For installation and upgrade steps, see Setting up Cloudera Iceberg REST Catalog for data sharing.

What's New in Apache Iceberg

Learn about the new features of Iceberg in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

Upgrade Hive Iceberg library to Apache Iceberg 1.10.0

In Cloudera Runtime 7.3.2 SP1, the Apache Iceberg library used by the Hive engine is upgraded to Apache Iceberg 1.10.0. Existing Iceberg V1 and V2 table operations are unchanged; you need not run a new SQL syntax for current workflows.

This upgrade aligns Hive with the platform-wide Iceberg 1.10.0 library and provides the foundation for Iceberg V3 capabilities in Hive. For supported Iceberg features, see the .

Apache Jira: HIVE-28665

Upgrade Impala Iceberg library to Apache Iceberg 1.10.0

In Cloudera Runtime 7.3.2 SP1, the Apache Iceberg library used by the Impala engine is upgraded to Apache Iceberg 1.10.0. Existing Iceberg V1 and V2 table operations are unchanged; you need not run a new SQL syntax for current workflows.

This upgrade aligns Impala with the platform-wide Iceberg 1.10.0 library and provides the foundation for Iceberg V3 capabilities in Impala. For supported Iceberg features, see the .

Apache Jira: IMPALA-14665

Hive Iceberg variant data type support

Hive supports the Iceberg VARIANT data type. The Apache Iceberg format specification defines this semi-structured type for Iceberg format version 3. You can store semi-structured data in VARI ANT columns and query nested fields using PARSE_JSON and VARIANT_GET (HIVE-29183).

For more information, see and .

Apache Jira: HIVE-29183

Hive Iceberg deletion vector support

Hive supports deletion vectors on Iceberg V3 tables. Deletion vectors are the Iceberg V3 encoding for position deletes and replaces the V2 position delete files for new delete operations on V3 tables. When you run DELETE on a V3 Iceberg table, Hive records deleted row positions using deletion vectors.

For more information, see and .

Apache Jira: HIVE-29006

Impala Iceberg deletion vector support [Technical Preview]



Note: This feature is in Technical Preview and is not ready for production deployments. Cloudera recommends trying this feature in test or development environments and encourages you to provide feedback on your experiences.

Impala supports deletion vectors on Iceberg V3 tables. Deletion vectors are the Iceberg V3 encoding for position deletes and replace V2 position delete files for new delete operations on V3 tables. Impala can read tables that contain deletion vectors written by Impala or by other Iceberg V3-compatible engines. When you run DELETE on a V3 Iceberg table from Impala, Impala records deleted row positions using deletion vectors.

For more information, see and .

Apache Jira: IMPALA-14585

Impala Iceberg row lineage support [Technical Preview]



Note: This feature is in Technical Preview and is not ready for production deployments. Cloudera recommends trying this feature in test or development environments and encourages you to provide feedback on your experiences.

Impala supports row lineage tracking on Iceberg format V3 tables. Row lineage is an Iceberg V3 capability that assigns a unique identifier to each row and records row provenance across table snapshots. This metadata supports change data capture (CDC) use cases. To use row lineage from Impala, create an Iceberg V3 table by setting 'format-version'='3'. When Impala writes to the table, Impala maintains the row lineage metadata required; you need not run a new SQL syntax to enable row lineage.

For more information, see .

Apache Jira: IMPALA-14592

Impala Iceberg puffin statistics support

Impala supports writing Puffin statistics files for Iceberg V3 tables. Puffin is a lightweight container format that stores binary metadata such as deletion vector sidecar files required by the Iceberg V3 specification.

For more information, see .

Apache Jira: IMPALA-14755

Impala Iceberg table repair support

Impala introduces the repair_metadata() function for Iceberg tables. This function provides a self-service recovery path to recover Iceberg tables that are inaccessible due to missing data files after manual file deletions in the underlying storage.

For more information, see .

Apache Jira: IMPALA-13725

Impala Iceberg LIKE predicate pushdown

Impala pushes eligible SQL LIKE predicates on Iceberg tables to Iceberg startsWith() and equa l() expressions during query planning. This enables Iceberg to use manifest statistics for file and partition pruning, which can reduce I/O on filtered string scans.

For more information, see .

Apache Jira: IMPALA-14737

Cloudera Lakehouse Optimizer

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Iceberg and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see New Features.

Cloudera Lakehouse Optimizer for Iceberg table optimization

In Cloudera Runtime 7.3.2 and higher versions, you can use Cloudera Lakehouse Optimizer service in Cloudera Manager to automate the Iceberg table maintenance tasks.

Cloudera Lakehouse Optimizer provides automated Iceberg table maintenance, through Spark jobs, for Iceberg tables in Cloudera Open Data Lakehouse. It simplifies table management, improves query performance, and reduces operational costs.

You can add the service to an existing Cloudera Base on premises 7.3.2 or higher versions cluster in Cloudera Manager 7.13.2 or higher versions, or you can create a dedicated cluster and then add the service. You must ensure that the cluster contains all the required services. After you finish configuring the service, you can use the Cloudera Lakehouse Optimizer service REST APIs to define the Cloudera Lakehouse Optimizer policies, perform policy management, and run other Iceberg table optimization operations.

For more information, see Cloudera Lakehouse Optimizer.

Integrate Iceberg scan metrics into Impala query profiles

Iceberg scan metrics are now integrated into the Frontend section of Impala query profiles, providing deeper insight into query planning performance for Iceberg tables.

The query profile now displays scan metrics from Iceberg's planFiles() API, including total planning time, counts of data/delete files and manifests, and the number of skipped files.

Metrics are displayed on a per-table basis. If a query scans multiple Iceberg tables, a separate metrics section will appear in the profile for each one.

Apache Jira: IMPALA-13628

Delete orphan files for Iceberg tables

You can now use the following syntax to remove orphan files for Iceberg tables:

```
-- Remove orphan files older than '2022-01-04 10:00:00'.
ALTER TABLE ice_tbl EXECUTE remove_orphan_files('2022-01-04 10:
00:00');

-- Remove orphan files older than 5 days from now.
```

```
ALTER TABLE ice_tbl EXECUTE remove_orphan_files(now() - interval
 5 days);
```

This feature removes all files from a table's data directory that are not linked from metadata files and that are older than the value of older_than parameter. Deleting orphan files from time to time is recommended to keep the size of a table's data directory under control.

Apache Jira: IMPALA-14492

Allow forced predicate pushdown to Iceberg

Since IMPALA-11591, Impala has optimized query planning by avoiding predicate pushdown to Iceberg unless it is strictly necessary. While this default behavior makes planning faster, it can miss opportunities to prune files early based on Iceberg's file-level statistics.

A new table property, impala.iceberg.push_down_hint is introduced, which allows you to force predicate pushdown for specific columns. The property accepts a comma-separated list of column names, for example, 'col_a, col_b'.

If a query contains a predicate on any column listed in this property, Impala will push that predicate down to Iceberg for evaluation during the planning phase.

Apache Jira: IMPALA-14123

UPDATE operations now skip rows that already have the desired value

The UPDATE statement for Iceberg and Kudu tables is optimized to reduce unnecessary writes.

Previously, an UPDATE operation would modify all rows matching the WHERE clause, even if those rows already contained the new value. For Iceberg tables, this resulted in writing unnecessary new data and delete records.

With this enhancement, Impala automatically adds an extra predicate to the UPDATE statement to exclude rows that already match the target value.

Apache Jira: IMPALA-12588

What's New in Apache Impala

New features and functional updates for Impala are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

New metrics for catalog server table invalidation

The catalog server now exposes metrics that report how many tables were invalidated due to time-to-live (TTL) expiration and memory pressure, including the rate of invalidations over several time windows, such as the past 10 seconds, 1 minute, 5 minutes, and 30 minutes. Administrators can use these metrics to tune settings such as invalidate_tables_timeout_s and the JVM heap size allocated to the catalog server.

Apache Jira: IMPALA-14684

Cloudera Runtime 7.3.2.10000 SP1

Plain-English query analysis

You can now get instant, plain-English performance summaries and clear root-cause recommendations directly inside the Impala Coordinator Web interface. The system securely analyzes query steps on demand by using an intelligent tool-calling framework to provide insights without hitting model data limits.

For more information, see .

Horizontal Data Aggregation with PIVOT

Cloudera Runtime SQL queries now support the `PIVOT` clause to transform unique column values into multiple new columns for horizontal data aggregation. This feature allows you to reorganize flat query output into a multi-dimensional grid report.

For more information, see the .

Data Normalization with UNPIVOT

Cloudera Runtime SQL queries now support the `UNPIVOT` clause to rotate multiple columns into rows within a single data column. This feature allows you to transform structured data sets into normalized results for aggregation and analysis.

For more information, see the .

Database Properties Support in Impala

Impala SQL queries now support setting database key and value properties by using the `WITH DBPROPERTIES` clause in `CREATE DATABASE` statements and the `SET DBPROPERTIES` clause in `ALTER DATABASE` statements.

For more information, see the .

User role management

You can now grant and revoke roles directly to and from individual users in Impala, providing more granular control over security management. This feature includes support for the `GRANT ROLE`, `REVOKE ROLE`, and `SHOW ROLE GRANT USER` statements, aligning Impala with Apache Hive's role-related functionality.

For more information, see the , and .

Apache Jira: IMPALA-14085

Cost-based query planning with the Apache Calcite planner [Technical Preview]

Impala can now compile queries with the Apache Calcite planner, a cost-based optimizer that evaluates alternative execution plans, such as different join orders, to choose a more CPU- and memory-efficient plan. The Calcite planner is disabled by default; you enable it per session by setting the PLANNER query option.

For more information, see .

Apache Jira: IMPALA-12871

Native geospatial query acceleration

Cloudera Runtime 7.3.2.10000 SP1 introduces native implementations for specific geospatial functions to accelerate simple queries. This feature reduces processing overhead by avoiding transitions to the Java Virtual Machine and optimizing file-level filtering for Parquet and Iceberg tables.

For more information, see .

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Impala and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see New Features.

Hierarchical metastore event processing (Preview)

Impala now supports a multi-layered, hierarchical approach to metastore event processing to improve synchronization speed and handle event dependencies more efficiently. By enabling this

feature, you can segregate events based on their dependencies and process them independently through a system of database and table event executors. This method reduces synchronization time for Hive Metastore (HMS) events by allowing parallel processing while maintaining linearizability for specific tables.

For more information, see

Impala AES encryption and decryption support

Impala now supports AES (Advanced Encryption Standard) encryption and decryption to work better with other systems. AES-GCM is the default mode for strong security, but you can also use other modes like CTR, CFB, and ECB for different needs. This feature works with both 128-bit and 256-bit keys and includes checks to keep your data safe and confidential.

For more information, see

Apache Jira: IMPALA-13039

Ubuntu 24.04 support

You can now build and run Impala on Ubuntu 24.04.

Dual-stack networking support

You can now configure Impala to support dual-stack networking, allowing the service to handle both IPv4 and IPv6 traffic simultaneously. This update includes new configuration properties in Cloudera Manager and support for dual-stack load balancing with HAProxy.

For more information, see Impala dual stack IPv6 and Impala dual stack HA Proxy

Default support for Java 17

Cloudera now provides Impala using Java 17 for builds and runtime environments.

OpenTelemetry integration for Impala

Cloudera now provides OpenTelemetry (OTel) support to help you see query performance and troubleshoot issues. This new feature, collects and exports query telemetry data as OpenTelemetry traces to a central OpenTelemetry compatible collector. The integration is designed to have a minimal impact on performance because it uses data already being collected and handles the export in a separate process. For more information, see

Apache Jira: IMPALA-13234

Caching intermediate query results

Cloudera now supports caching intermediate results to improve query performance and resource efficiency for repetitive workloads. By storing results at various locations within the SQL plan tree, the system can reuse computation for similar queries even when they are not identical, provided the underlying data and settings remain unchanged. For more information, see

Impala-shell and Impyla now supports Python 3.12

Cloudera now provides support for Python 3.12 in `impala-shell` and Impyla.

Apache Jira: IMPALA-14452

Query cancellation supported during analysis and planning

This new feature allows you to cancel Impala queries even while they are in the Frontend stage, which includes analysis and planning. Previously, you could not cancel a query while it was waiting for operations like loading metadata from the Catalog Server. With this update, Impala now registers the planning process and can interrupt it to cancel the query.

Apache Jira: IMPALA-915

Improved memory estimation and control for large queries

Impala now uses a more realistic approach to memory estimation for large operations like SORT, AGGREGATION, and HASH JOIN.

Previously, these operations could severely overestimate their memory needs (sometimes requesting terabytes) when row counts were misestimated. This often caused the Admission Controller to reject the query outright, even though the operation could easily handle the data by writing (spilling) to disk.

The system now considers the operator's ability to spill data to disk and caps the memory estimate based on your cluster's actual memory limits (like MEM_LIMIT). This change makes more large queries admittable and allows them to run successfully.

New Control Option: A new query option, MEM_ESTIMATE_SCALE_FOR_SPILLING_OPERATOR (a scale from 0.0 (exclusive) to 1.0 (inclusive)), is introduced, giving you control over this behavior:

- Higher values (closer to 1.0): Tells Impala to reserve more memory, increasing the chance of faster, in-memory execution.
- Lower values (closer to 0.0, but NOT 0.0): Tells Impala to request less memory, increasing the chance your query is admitted, but potentially leading to more spilling to disk (slower execution).
- The default value is 0.0 which disables this feature, and revert Impala planner to old behavior.

Apache Jira: IMPALA-13333

Expose query cancellation status to UDF interface

Impala now exposes the query cancellation status to the User-Defined Function (UDF) interface. This new feature allows complex or time-consuming UDFs to periodically check if the query has been cancelled by the user. If cancellation is detected, the UDF can stop its work and fail fast.

This significantly reduces the time you have to wait to stop a long-running query that is stuck inside a UDF evaluation.

Apache Jira: IMPALA-13566

Expanded compression levels for ZSTD, and ZLIB

Impala has extended the configurable range of compression levels for ZSTD, and ZLIB (GZIP/DEFLATE) codecs. This enhancement allows for better optimization of the trade-off between compression ratio and write throughput.

- ZSTD: Supports a wider range, including negative levels, up to 20.
- ZLIB (GZIP, DEFLATE): Supports levels from 1 (default) to 9 (best compression).

These levels are applied via the compression_codec query option.

JIRA Issue: IMPALA-13923

Constant folding is now supported for non-ASCII and binary strings

Previously, the query planner could not apply the optimization known as constant folding if the resulting value contained non-ASCII characters or was a non-UTF8 binary string. This failure meant that important query filters could not be simplified, which prevented key performance optimizations like predicate pushdown to the storage engine (e.g., Iceberg or Parquet stat filtering).

The planner is updated to correctly handle and fold expressions resulting in valid UTF-8 strings (including international characters) and binary byte arrays. This allows Impala to push down more filters, significantly improving the performance of queries that use non-ASCII string literals or binary data in their filters.

JIRA Issue: IMPALA-10349

Catalogd and Event Processor Improvements

- Faster Inserts for Partitioned Tables (IMPALA-14051): Inserting data into very large partitioned tables is now much faster. Previously, Impala communicated with the Hive Metastore (HMS) one partition at a time, which was a major slowdown. Impala now uses the batch insert API to

send all insert information to the HMS in one highly efficient call, significantly boosting the performance of your `INSERT` statements into transactional tables.

- Quicker Table Administration (IMPALA-13599): Administrative tasks, such as running `DROP STATS` or changing the CACHED status of a table, are now much faster on tables with many partitions. Impala previously made thousands of individual calls to the HMS for these operations. The system now batches these updates, making far fewer calls to the HMS and speeding up these essential administrative commands.
- Reliable Table Renames (IMPALA-13989): The `ALTER TABLE RENAME` command no longer fails when an `INVALIDATE METADATA` command runs at the same time. Previously, this caused the rename to succeed in the Hive Metastore but fail in Impala's Catalog Server. Impala now includes automatic error handling that instantly runs an internal metadata refresh if the rename is interrupted, ensuring the rename completes successfully without requiring any manual user steps.
- Efficient Partition Refreshes (IMPALA-13453): Running `REFRESH <table> PARTITION <partition>` is now much more efficient. Previously, this command always fully reloaded the partition's metadata and column statistics, even if the partition was unchanged. Impala now checks if the partition data has changed before reloading, avoiding the unnecessary drop-add sequence and significantly improving the efficiency of partition metadata updates.
- Reduced Partition API Calls (IMPALA-13599): Impala has reduced unnecessary API interactions with the HMS during table-level operations. Commands like `ALTER TABLE... SET CACHED/UNCACHED` or `DROP STATS` on large tables previously generated thousands of single alter_partition() calls. Impala now utilizes the HMS's bulk-update functionality, batching these partition updates to drastically reduce the total number of required API calls.
- `REFRESH` on multiple partitions (IMPALA-14089): Impala now supports using the `REFRESH` statement on multiple partitions within a single command, which significantly speeds up metadata updates by processing partitions in parallel, reduces lock contention in the Catalog service, and avoids unnecessary increases to the table version. See
- Impala cluster responsiveness during table renames(IMPALA-13631): This ensurs that the critical internal lock is no longer held during long-running external calls initiated by `ALTER TABLE RENAME` operations. This prevents the entire Impala cluster from being blocked, allowing other queries and catalog operations to proceed without interruption.

Apache Jira: IMPALA-14051, IMPALA-13599, IMPALA-13989, IMPALA-13453,IMPALA-14089 , IMPALA-13631

Auto-optimized parquet collection queries

Impala now automatically boosts query performance for tables with collection data types by setting the parquet_late_materialization_threshold to 1 when data can be skipped during filtering. This ensures maximum efficiency by reading only the data needed.

For more information, see

Apache Jira: IMPALA-3841

Impala-shell now shows row count and elapsed time for most statements in HiveServer2 mode

When running Impala queries, some commands over HiveServer2 protocol (like REFRESH or INVALIDATE) did not show the Fetched X row(s) in Ys output in Impala-shell, even though Beeswax protocol showed them.

This issue was resolved by adding a new option in Impala-shell called --beeswax_compat_num _rows. When this option is enabled, Impala-shell now prints Fetched 0 row(s) in along with the elapsed time for all Impala commands. This requires impala-shell 4.5.1a1 or higher. See,

Apache Jira: IMPALA-13584

Support for arbitrary encodings in text and sequence files

Impala now supports reading from and writing to Text and Sequence files that use arbitrary character encodings, such as GBK, beyond the default UTF-8.

Impala now recognizes the Hive table property "serialization.encoding" used with LazySimpleSe rDe.

- When reading data from these files, Impala uses the specified encoding to correctly decode the string data (example, converting GBK bytes into readable characters).
- When inserting data into these tables, Impala correctly encodes the inserted strings into the specified charset before saving them to the text file.

For more information, see .

Apache JIra: IMPALA-10319

New query options for reliable metadata synchronization

Impala now offers new query options to give you a reliable way to ensure your queries run with the latest table data after the relative HMS modifications are done.

- SYNC_HMS_EVENTS_WAIT_TIME_S: Sets the maximum time, in seconds, you are willing to wait for the metadata synchronization from the Hive Metastore (HMS). Setting this enables Impala to pause query compilation automatically until changes are applied, ensuring metadata consistency.
- SYNC_HMS_EVENTS_STRICT_MODE: Controls error handling if the wait time is exceeded. By default, Impala proceeds with a warning. Set to TRUE to force the query to fail immediately, guaranteeing strict consistency.

See,

JIRA Issue: IMPALA-12152

Impala now supports Hive’s legacy timestamp conversion to ensure consistent interpretation of historical timestamps

When reading Parquet or Avro files written by Hive using legacy timestamp conversion, Impala's timezone calculation for UTC timestamps could be incorrect, particularly for historical dates and timezones like Asia/Kuala_Lumpur or Singapore before 1982. This meant the timestamps displayed in Impala were different from those in Hive.

This issue is addressed by Impala now checks for the writer.zone.conversion.legacy flag in the Parquet file metadata to determine if Hive's legacy timestamp conversion method was used. If found, Impala uses a compatible conversion method. For older files without this flag, a new session and query option, use_legacy_hive_timestamp_conversion, has been added to control the conversion method. See,

Apache Jira: IMPALA-13627

Supporting one-dimensional arrays in Kudu tables

Impala now supports one-dimensional arrays in Kudu tables. You can create Kudu tables with array columns and perform selection queries on these complex collection types.

Impala now supports tables created with the Hive JDBC Storage handler

Impala now translates and adapts the hive JDBC table properties upon loading, making them compatible with Impala's internal requirements. Previously, Impala had difficulty reading tables created using the hive JDBC Storage handler due to differences in how table properties, such as JDBC driver and DBCP configurations, were defined compared to Impala-created tables. See,

Apache Jira: IMPALA-12992

New catalogd flag to disable HMS sync by default

You can now use the disable_hms_sync_by_default `catalogd` startup flag to set a global default for the impala.disableHmsSync property. This feature allows you to skip event processing for all databases and tables by default while opting in specific elements as needed.

For more information, see:

Apache Jira: IMPALA-14131

Parallel metadata loading in local catalog mode

Previously, when a query accessed multiple unloaded tables in local catalog mode, Impala loaded the metadata for those tables one after another. This sequential process caused significant latency and performance regressions compared to the legacy catalog mode.

This issue is addressed by parallelizing the table loading process. The fix allows Impala to load and gather metadata for multiple tables simultaneously. You can control the maximum number of threads used for this process by using the new max_stmt_metadata_loader_threads flag, which defaults to 8 threads per query compilation. See,

Apache Jira: IMPALA-14447

Specifying compression levels for LZ4, ZLIB, and ZSTD

You can now specify compression levels for the LZ4, ZLIB, GZIP, and ZSTD codecs to achieve higher compression ratios. This includes support for high compression modes in LZ4 (levels 3–12) and negative compression levels for ZSTD. These levels are supported by using the compression_codec query option.

For more information, see

Apache Jira: IMPALA-10630, IMPALA-14082

Impala now supports ARM architecture

What's New in Apache Kafka

New features and functional updates for Kafka are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

Centralized TLS protocol and cipher settings

When you enable the TLS_ADVANCED_CONTROL feature flag, Kafka Broker, KRaft Controller, Kafka Connect, and MirrorMaker 1 roles expose Supported SSL/TLS versions (ssl.enabled.protocols) and TLS Cipher List (ssl.cipher.suites) properties in Cloudera Manager. These properties default to the centralized Cloudera Manager TLS settings. For information about enabling the feature flag, see *Transport Layer Security Support and Custom Protocol Configuration*.

Rebase on Apache Kafka 3.9.2

Kafka shipped with Cloudera Runtime 7.3.2.10000 SP1 is based on Apache Kafka 3.9.2 (previously 3.9.1). This maintenance rebase includes critical bug fixes and security updates.

For the complete list of upstream changes, see the Apache Kafka 3.9.2 release notes. Notable upstream fix: KIP-1252 (AlterConfigPolicy behavior in KRaft vs ZooKeeper).

KRaft Controller auto joining

The KRaft Controller auto joining feature from upstream Apache Kafka 4.2 is available in Cloudera Runtime 7.3.2 SP1. The controller.quorum.auto.join.enable configuration currently has true as the default value that enables seamless KRaft Controller scaling. When this configuration is enabled, a KRaft Controller role tries to join the KRaft Quorum as soon as the role starts. KIP-853 contains more information about the feature. See KIP-853: KRaft Controller Membership Changes.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Kafka and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see New Features.

Rebase on Kafka 3.9

Kafka shipped with this version of Cloudera Runtime is based on Apache Kafka 3.9.1 (previously 3.4.1). For more information, see the following resources:

- Notable changes for releases 3.5.0 through 3.9.1: Upgrading | Apache Kafka.
- The Apache Kafka release notes for the following versions:
  - Kafka 3.5.0
  - Kafka 3.5.1
  - Kafka 3.6.0
  - Kafka 3.6.1
  - Kafka 3.5.2
  - Kafka 3.7.0
  - Kafka 3.6.2
  - Kafka 3.7.1
  - Kafka 3.8.0
  - Kafka 3.8.1
  - Kafka 3.9.0
  - Kafka 3.9.1
- The Apache Kafka release announcements: Release Announcements | Apache Kafka

KRaft is generally available and ZooKeeper is deprecated

KRaft (Kafka Raft) is generally available. KRaft is from now on the recommended metadata management mode for Kafka in Cloudera. Additionally, migrating existing ZooKeeper-based Kafka clusters to use KRaft is now possible.

With the general availability of KRaft, deploying new or using existing Kafka clusters running in ZooKeeper mode is deprecated. Additionally, support for ZooKeeper-based Kafka clusters will be removed in a future release.

Cloudera recommends the following:

- Deploy all new Kafka clusters in KRaft mode.
- Migrate existing ZooKeeper-based clusters to KRaft following an upgrade to Cloudera Runtime 7.3.2.

  This is the only version where migration is possible. Neither previous or future major, minor, and maintenance versions support migration.

For additional information, see the following resources:

- Installing Kafka in Cloudera Base on premises
- Kafka KRaft
- Migrating Kafka from ZooKeeper to KRaft overview

Kafka protocol and metadata version is set automatically during upgrades

When upgrading Kafka, Cloudera Manager now automatically sets the inter.broker.protocol.version property for ZooKeeper-based clusters and the metadata.version property for KRaft-based clusters. You no longer need to manually set these properties to the current protocol or metadata version before an upgrade. This feature is only available when upgrading to Cloudera Runtime 7.3.2 or higher.

After the upgrade, clearing these properties remains a manual task. However, in Cloudera Runtime 7.3.2 and higher, both inter.broker.protocol.version and metadata.version are now available for

direct configuration in Cloudera Manager Kafka Configuration . The label names of the properties are Kafka Inter-Broker Protocol Version and Kafka Metadata Version. This means you can set or clear these properties directly from the UI, without needing to use advanced configuration snippets.

Connector-level offset flush control

A new connector-level property, cloudera.offset.flush.interval.ms, is added. Use this property to override the Kafka Connect role-level Offset Flush Interval (offset.flush.interval.ms) property. Overriding enables you to control the interval at which connector task offsets are committed on a per-connector basis.

Configure cloudera.offset.flush.interval.ms in connectors that need a different offset flush interval than the role default. This is commonly useful for connectors where the interval controls how often data is flushed to target systems, for example NiFiStatelessSink, HDFSSink, and S3Sink.

IPv6 support for Kafka

Starting with the 7.3.2 release, Kafka supports IPv6 with dual-stack functionality, allowing seamless communication over both IPv4 and IPv6 networks. This capability improves network scalability, future-proofs deployments, and enhances overall platform security.

Offline Log Directories chart

A new default chart, **Offline Log Directories**, is added for Kafka in Cloudera Manager. This chart can help you quickly identify and track storage issues on your brokers. It is available by default for the Kafka service as well as for individual Kafka Broker role instances.

The chart shows offline log directories and their mount paths for Kafka brokers. A non-zero value indicates an active error state for a specific log directory, while a value of 0 means the directory was in an error state during the selected timeframe but is now healthy. The chart only displays log directories that had errors during the selected timeframe.

New actions for collecting Kafka diagnostic data

The following new service-specific actions are available for collecting Kafka diagnostic data in Cloudera Manager:

- Collect Kafka Cluster Diagnostics - gathers detailed cluster-wide data, including topics, configurations, consumer groups, and more.
- Describe Kafka Topics - provides detailed information about all Kafka topics.

These actions are available in the Actions dropdown on the Kafka service and Kafka Broker role instance pages. Diagnostic data is printed to sdtout for immediate access and also saved as a compressed archive on the host where the action runs.

For more information, seeCollecting Kafka diagnostic data using Cloudera Manager actions.

Debezium connectors upgraded from 1.9.8.Final to 3.3.1.Final

This release of Cloudera Runtime ships version 3.3.1.Final of the following Debezium connectors:

- MySQL
- PostgreSQL
- Oracle
- SQL Server
- Db2

Existing connector instances are automatically upgraded to the new version as part of a cluster upgrade. However, you will be required to make configuration updates before you can upgrade your cluster. Critical changes that affect all Debezium connectors are summarized below.



Important: Debezium 3.3.1.Final includes many major breaking changes compared to 1.9.8.Final. Depending on your connector configuration and use case, additional configuration updates not highlighted here might be necessary. For a comprehensive list of changes, review release notes on https://debezium.io/.

- Property renaming (configuration namespace changes)

  New, more consistent namespaces for configuration properties are introduced. The old database.* prefixes have been removed. Connector configuration keys collected in the following table must be updated before an upgrade.

Old Property Prefix (Debezium 1.9)	New Property Prefix (Debezium 3.3)
database.server.name	topic.prefix
database.history.*	schema.history.internal.*
database.* (JDBC pass-through)	driver.*
database.dbname (SQL Server)	database.names

- Database driver version requirements are updated

  The recommended and supported JDBC driver versions used by the majority of connectors has changed. The following table collects the JDBC drivers you will need to deploy on your cluster before an upgrade.

  Note: The recommended driver version for Debezium Db2 remains unchanged and is still 11.5.0.0 but is listed in the following table as a reference.

Component	New Driver Version / Notes
MySQL	9.1.0
PostgreSQL	42.7.7
Oracle	21.x, 23.x — use a Java 11+ Oracle JDBC driver (ojdbc11.jar)
SQL Server	12.4.2.jre8
Db2	11.5.0.0

For step-by-step upgrade instructions on configuration updates and driver replacements, see Getting Started Upgrading a Cluster.

What's New in Apache Knox

New features and functional updates for Apache Knox are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2:

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1:

Configurable TLS settings for Knox service discovery

By default, Knox service discovery uses the global TLS cipher and protocol settings when communicating with Cloudera Manager. When Cloudera Manager Server TLS settings differ from the global cluster configuration, you can now configure matching TLS settings on both Cloudera Manager Server and Knox service discovery. You can configure gateway.cloudera.manager.service .discovery.ssl.protocols and gateway.cloudera.manager.service.discovery.ssl.ciphers in Cloudera Manager for Knox service discovery.

For more information about Knox service discovery TLS settings, see Configuring TLS/SSL encryption manually for Apache Knox. For Cloudera Manager Server TLS configuration, see Manually Configuring TLS Encryption for Cloudera Manager.

Global TLS configuration for Apache Knox

Starting in Cloudera Runtime 7.3.2 SP1 (7.3.2.10000), Knox supports Cloudera Manager global TLS protocol and cipher settings. When you enable the CMF_FF_TLS_ADVANCED_CONTROL feature flag, you must configure TLS through the Supported SSL/TLS versions and TLS Cipher List parameters on the Knox service configuration page. When the feature flag is disabled, you must configure TLS through the Knox TLS - Protocols and Knox TLS - Cipher Suites parameters on the Knox service configuration page.

For more information, see Configuring TLS/SSL encryption manually for Apache Knox.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Knox and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see New Features.

SameSite attribute for pac4j session cookies is now configurable

You can now configure the SameSite attribute for pac4j session cookies.

Group impersonation support in Knox

Knox now supports group impersonation, allowing users in specific groups to impersonate other users. For more information, see Configuring Group Impersonation in Knox.

Knox IDBroker integration with HashiCorp Vault

Knox IDBroker now integrates with HashiCorp Vault for AWS credentials management, allowing IDBroker to authenticate with AWS using short-lived credentials from Vault instead of storing long-lived credentials for this purpose. For more information, see Configuring Knox IDBroker with HashiCorp Vault.

Role-level alias management for Knox Gateway and IDBroker

The alias management configuration has been moved from service-level to role-level. Each role now has its own dedicated configuration: gateway_save_alias_command_input for the Knox Gateway role and idbroker_save_alias_command_input for the IDBroker role. Two role-specific commands are now available: Save Alias - Knox Gateway and Save Alias - IDBroker. For more information, see Saving aliases.

What's New in Apache Kudu

New features and functional updates for Kudu are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

Prometheus metrics integration with Grafana

Kudu provides native integration with Prometheus by exposing metric endpoints on Master and Tablet Server daemons for cluster monitoring. Prometheus historical retention, dynamic service discovery, and Grafana visualization allow you to track performance, configure alerts, and build custom dashboards across nodes.

For more information, see Prometheus metrics integration with Grafana .

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Kudu and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see New Features.

Kudu JDK 17 upgrade

Kudu supports JDK 17. Upgrading to JDK 17 ensures your application stays competitive and maintainable by providing improved performance, increased security, modern language features, and long-term support.

IPv6 support for Kudu

Starting with the 7.3.2 release, Kudu supports IPv6 with dual-stack functionality, allowing seamless communication over both IPv4 and IPv6 networks. This capability improves network scalability, future-proofs deployments, and enhances overall platform security.

For more information, see Enable Dual-Stack support for Kudu.

Support for one-dimensional arrays in Kudu

Kudu now supports one-dimensional arrays for scalar types, such as integers, strings, booleans, and binaries. This feature allows you to store multiple related values in a single column, which improves query performance and simplifies your data model by removing the need for complex joins or denormalized tables.

For more information, see Manage one-dimensional arrays in Kudu.

Replicating Kudu tables using Apache Flink

You can now use the Apache Flink-based replication job to continuously replicate data from a source Kudu cluster to a destination Kudu table. This feature uses Kudu diff scans to efficiently capture and move only the changed rows, supporting high-availability and disaster recovery scenarios.

For more information, see Replicating Kudu tables using Apache Flink.

What's New in Livy

New features and functional updates for Livy are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

FIPS 140-3 support with Spark 3

Livy for Spark 3 supports FIPS 140-3 on FIPS-enabled clusters alongside Spark 3 in this service pack.

For configuration details, see .

Encryption in transit for Livy Server

Livy for Spark 3 inherits shared Spark 3 configuration from Cloudera Manager and supports TLS for the Livy Server REST API, including TLS 1.3 when your Cloudera Manager TLS policy allows it.

For configuration details, see , , and .

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2.0

Livy for Spark 3 in Cloudera Runtime

Livy for Spark 3 is the Livy service in Cloudera Runtime 7.3.2.0. It submits and manages Spark 3 sessions using the default Apache Spark 3.5.4 runtime shipped with this release.



Important:

Apache Livy 2 is not available in Cloudera Runtime 7.3.2. Livy for Spark 3 is the supported Livy service.

Refer to Upgrading Spark for more information on upgrading Spark clusters to 7.3.2.0, and Migrating Spark applications for more information on migrating your existing Spark applications between versions 2 and 3.

What's New in Navigator Encrypt

Learn about the new features of Navigator Encrypt in Cloudera Runtime 7.3.2.0, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

TLS 1.3 support

Navigator Encrypt supports TLS 1.3 for stronger security and compliance.

FIPS support

Navigator Encrypt is now supported in FIPS mode.

Support for SLES15 SP5, SLES15 SP7, and Ubuntu 24

Navigator Encrypt supports SLES15 SP5, SLES15 SP7, and Ubuntu 24.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features in this release.

Cloudera Runtime 7.3.2.0

Navigator Encrypt installed using Cloudera Manager

In Cloudera Runtime 7.3.2.0 and higher releases, the new Navigator Encrypt parcel centralizes the entire lifecycle management within Cloudera Manager. Previously, Navigator Encrypt required manual, host-by-host RPM installation and configuration. Administrators can now manage the installation, configuration, and upgrade of Navigator Encrypt across all cluster hosts directly from the Cloudera Manager interface. This new approach provides a significant improvement in usability, streamlining ongoing maintenance and lifecycle management.

For more details, see Navigator Encrypt installed using Cloudera Manager.

What's New in Oozie

New features and functional updates for Oozie are introduced in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new features in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features in this release.

Cloudera Runtime 7.3.2

Configurable Oozie SSH action port

You can configure the port number for Oozie SSH actions. While the default port number remains 22, you can override this value using the following methods:

1. On workflow level, use the new 0.4 schema version for your SSH action in your workflow.xml file and add the port XML element with the new value.

   Example

   ```
   <port>11100</port>
   ```

2. On global level, add the oozie.action.ssh.action.port property to the Oozie Server Advanced Configuration Snippet (Safety Valve) for oozie-site.xml advanced configuration snippet in Cloudera Manager.

   Example

   ```
   "oozie.action.ssh.action.port: 11100"
   ```

If you set both options, then the value defined in the workflow.xml file takes precedence. Ensure that the port with the new port number opens in the target host.

What's New in Ozone

New features and functional updates for Ozone are introduced in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

Ozone S3 Security Token Service (STS)

Ozone introduces the S3 Security Token Service (STS), enabling administrators to issue short-lived, temporary S3 credentials using an AWS-compatible AssumeRole API. This is designed for data-lake workloads and applications that require temporary, scoped access to Ozone buckets and keys without distributing long-lived credentials tied to Kerberos identities.

Key features and capabilities:

- **AWS-compatible AssumeRole API**: Programmatically exchange permanent credentials for temporary ones.

- **Dedicated STS endpoint**: Runs on a separate endpoint on the S3 Gateway (isolated from the standard S3 object API port).
- **Ranger integration**: Authorization is handled using Ranger roles and policies. Users must be granted assume_role permissions on specific roles.
- **Inline session policies**: Optionally narrow permissions further using inline IAM session policies (which are intersected with the role's permissions).
- **Temporary credentials structure**:
  - AccessKeyId: Begins with ASIA.
  - SecretAccessKey: Standard secret key.
  - SessionToken: Opaque, server-side stateless token that must be passed in subsequent requests via the x-amz-security-token header.
- **Configurable token lifetimes**: Temporary credentials remain valid anywhere from 15 minutes to 12 hours (900 to 43,200 seconds), defaulting to 1 hour if unspecified.

For more information, see Ozone S3 Security Token Service (STS).

Cloudera Runtime 7.3.2.10000 SP1

Ozone Disk Balancer

Ozone includes an intra-node Disk Balancer tool that automatically detects storage imbalances and redistributes data evenly across all disks within a single DataNode.

This feature provides the following key capabilities:

- **Intra-node load balancing**: Eliminates I/O hotspots and performance bottlenecks by balancing utilization across individual local volumes.
- **Automated skew mitigation**: Resolves disk space imbalances caused by adding new disks, replacing failed drives, or processing bulk block deletions.

For more information, see Disk Balancer.

Ozone Cluster Capacity view in Recon UI

The Cluster Capacity dashboard in Ozone Recon provides a centralized view of storage utilization across your entire cluster. Moving from high-level physical capacity to individual DataNode diagnostics, it tracks usable space, logical data distribution, pre-allocated storage, and pending deletion lifecycles to help you monitor storage health and balance capacity effectively. For more information, see Cluster Capacity.

HttpFS proxy user support in Ozone

Ozone HttpFS Gateway supports proxy user impersonation, enabling authorized service accounts such as Hive and Data Explorer to perform operations on behalf of other users. For more information, see Configuring HttpFS proxy users.

Configurable replica counts for diverged quasi-closed containers

Ozone now features configurable replica counts for diverged quasi-closed containers, enhancing data resilience.

Previously, Ozone maintained a fixed two copies for each origin of a diverged quasi-closed container. With this update, you can customize replica counts to allocate more copies to the origin with the highest Block Commit Sequence ID (BCSID). Increasing the replication factor for the healthiest, most up-to-date origin significantly reduces the risk of data loss during node failures.

The following new dynamic configuration keys are now available to manage these replica counts:

Table 6: Replica count configuration keys

Configuration key	Description	Default value
hdds.scm.replication.quasi.closed.stuck.best.origin.copies	Defines the number of copies to maintain for the origin with the highest BCSID among healthy replicas.	3
hdds.scm.replication.quasi.closed.stuck.other.origin.copies	Defines the number of copies to maintain for all other origins of the same container.	2

Enhanced container health monitoring in Storage Container Manager (SCM)

Ozone Storage Container Manager (SCM) includes a new ContainerHealthState enumeration that tracks and captures container replication statuses within container information. This enhancement enables efficient querying and monitoring of container health conditions and provides improved visibility into the replication state and overall health of containers across the cluster.

Recon container health monitoring enhancements

Ozone Recon includes an improved container health monitoring mechanism through Containe rHealthTaskV2. This update ensures that Recon uses the SCM Replication Manager as the single authoritative source of truth for container health status.

Key improvements include:

- **Improved data consistency**: Eliminates discrepancies between Recon and SCM by aligning Recon's health view with SCM's authoritative state.
- **Enhanced tracking**: Ensures containers known to SCM but not yet synchronized to Recon are accurately tracked for health issues.

Ozone Recon OM synchronization metrics for staging and queue-based processing

Ozone now includes metrics at various stages of the staging and queue-based events processing pipeline for Recon OM (Ozone Manager) synchronization. These metrics provide fine-grained observability into the Recon OM synchronization workflow, enabling operators to monitor the health and performance of the staging DB-based reprocessing of the OM DB.

Improved metadata scalability for Recon tasks

Ozone Recon now stores metadata for specific Ozone Manager (OM) background tasks in RocksDB instead of SQL Derby. This staging-based architecture improves throughput, scalability, and resource efficiency during heavy metadata synchronization.

Key Improvements include:

- **Enhanced performance**: Reduces overhead and locking issues during concurrent metadata processing.
- **Better scalability**: Handles large-scale datasets for file size distributions and table insights without SQL database bottlenecks.
- **Targeted metadata migration**: Applies specifically to file size counts (FILE_COUNT_BY_SIZE) and global table statistics (GLOBAL_STATS).

Accelerated replicated disk usage calculations in Recon

Recon optimizes replicated disk usage calculations by introducing a dedicated replicatedSizeOfFiles field in the Namespace Summary (NSSummary), eliminating recursive directory tree traversals.

Key Improvements include:

- **Faster UI performance**: Accelerates page loading times and improves responsiveness when viewing disk usage metrics in the Recon UI.
- **Cluster scalability**: Prevents performance bottlenecks in large clusters by eliminating heavy, recursive lookups.
- **Accurate capacity tracking**: Delivers precise storage metrics, including replication overhead, for effective cluster management and capacity planning.

Enhanced storage capacity visibility in DataNodes

The Storage Capacity table on the Ozone Recon DataNodes page now includes File System (FS) Reserved Space alongside existing capacity metrics.

Key Improvements include:

- **Accurate capacity tracking**: Displays reserved disk space directly within storage metrics to distinguish system-reserved space from total usable capacity.
- **Improved cluster visibility**: Helps you monitor actual available storage and prevent unexpected capacity constraints across DataNodes.

Ozone supports configurable RocksDB bottommost level compaction strategy

You can now configure the BottommostLevelCompaction strategy for RocksDB in Ozone. Previously hardcoded to kForce, this setting is now adjustable through configuration and CLI commands. This allows for better optimization of background compaction processes by choosing between kSkip, kForce, or kIfHaveCompactionFilter based on your cluster's performance needs.

Ozone volume health check toggle

Ozone introduces a configuration switch that allows you to enable or disable the RocksDB checks during volume health checks. This provides administrators with the ability to customize health check validations based on their specific environment requirements.

Added new container-to-key mapping tool

A new debugging tool allows you to quickly map containers to their associated keys. Input a list of containers to retrieve all matching keys across different layout types:

- FSO (File System Optimized)
- OBS (Object Store)
- Open (Uncommitted keys)
- MPU (Multipart Upload keys)

This simplifies key tracking and makes troubleshooting container data faster and more efficient.

Ozone Storage Container Manager (SCM) Replication Manager supports configurable container sample limits

Ozone SCM Replication Manager allows administrators to dynamically configure the number of containers sampled in each iteration. Previously, SCM hardcoded this limit to 100 containers per state per iteration. This improvement provides greater flexibility for monitoring container health and managing cluster state in large-scale deployments without requiring a service restart.

This enhancement introduces the following configuration property:

- **Property name**: hdds.scm.replication.container.sample.limit
- **Default value**: 100
- **Description**: Defines the maximum number of containers in each state to sample per Replication Manager iteration.
- **Dynamic update**: Supported. You can apply changes dynamically at runtime using the `ozone admin reconfig` command without restarting SCM.

NonWritableNodes metric added to Storage Container Manager

The Storage Container Manager (SCM) introduces the NonWritableNodes metric, providing administrators with a consolidated count of DataNodes currently unable to accept new write operations.

A DataNode is included in this count if it meets any of the following conditions:

- The node is not simultaneously in both HEALTHY and IN_SERVICE states (for example, nodes in ENTERING_MAINTENANCE or IN_MAINTENANCE states).
- The node has reached its maximum container allocation limit.
- The node is unable to write data to existing containers.

This metric improves cluster observability by replacing complex checks with a single data point for monitoring write availability across the cluster.

Support for deleting empty QUASI_CLOSED containers

You can now delete QUASI_CLOSED containers when all available replicas are completely empty, regardless of their current lifecycle state. To protect against data loss from lagging nodes, if a previously unseen or offline replica later reconnects with a higher Block Container Sequence ID (BCSID), the container is automatically restored to its original state. This enhancement enables clusters to safely reclaim metadata and storage space from empty, unclosed containers while guaranteeing that the system can still recover and re-replicate data if a newer replica emerges.

Ozone Container Balancer expands support for non-standard containers

The Ozone Container Balancer can now move containers that contain non-closed replicas (such as QUASI_CLOSED), provided the container meets minimum closed replica requirements. Previously, the balancer only processed containers where all replicas were strictly in a CLOSED state. This improvement optimizes data distribution and cluster balance in scenarios where certain replicas have not fully transitioned to a closed state despite the container being finalized in Storage Container Manager (SCM).

This feature is disabled by default. To enable it, set hdds.container.balancer.include.non.standard.containers configuration to true.

When enabled, the Container Balancer can process the following container states:

- CLOSED and OVER_REPLICATED containers with CLOSED replicas: Includes containers that meet the minimum required number of closed replicas based on their replication configuration (such as Ratis or Erasure Coding schemes).
- CLOSED and OVER_REPLICATED containers with QUASI_CLOSED replicas: Includes containers with QUASI_CLOSED replicas, provided the minimum required closed replicas are present.
- QUASI_CLOSED and OVER_REPLICATED containers: Includes containers where all replicas are QUASI_CLOSED and the source replica is non-empty.
- QUASI_CLOSED and HEALTHY containers: Includes containers where all replicas are QUASI_CLOSED and the source replica is non-empty.

The `ozone container info` command now displays Block Commitment Sequence ID (BCSID)

The `ozone container info` command now includes the BCSID directly in its standard console output alongside replica details. Previously accessible only in detailed JSON output, surfacing the BCSID in the default summary view improves container observability, enabling administrators to quickly verify container states and troubleshoot replica consistency without parsing complex JSON data.

Volume health counters added to `ozone admin datanode list` command

The `ozone admin datanode list` command now displays total and healthy volume counters for each DataNode directly in the cluster-wide output. By surfacing a clear healthy-to-total volume ratio (for example, 11/12), this enhancement improves cluster observability, enabling administrators to quickly identify disk degradation across the cluster and run detailed diagnostic queries like `ozone admin datanode info <datanode-id>` only when a volume discrepancy is found.

New latency breakdown metrics for key creation

Ozone now provides granular latency breakdown metrics for key creation operations, capturing execution time across each major stage of the creation process. By tracking individual sub-operations, these new metrics improve cluster observability, enabling administrators to quickly isolate sources of delay, pinpoint performance bottlenecks, and troubleshoot storage layer slowdowns.

Ozone SCM safemode status command supports checking all SCM nodes

Ozone enhances Storage Container Manager (SCM) safemode monitoring by adding the --all flag to the `ozone admin safemode status` command. This update allows you to view the safemode status of all SCM nodes in a High Availability (HA) setup simultaneously, making it easier to verify SCM state during startup and maintenance windows.

Parallel task processing for Ozone Recon

Ozone Recon now supports parallel processing for Ozone Manager (OM) RocksDB table iterations, significantly reducing database rebuild and bootstrap times on large clusters. Background operations including container-key mapping, file-size counting, and table insights now execute in parallel.

Administrators can fine-tune parallelism, thread allocation, and memory usage using four new configuration properties:

Property	Default value	Description
ozone.recon.task.reprocess.max.iterators	5	Sets the maximum number of iterators used during task reprocessing.
ozone.recon.task.reprocess.max.workers	20	Sets the maximum number of worker threads for parallel processing.
ozone.recon.task.reprocess.max.keys.in.memory	2000	Sets the maximum number of keys held in memory during reprocessing.
ozone.recon.filesizecount.flush.db.max.threshold	200000	Sets the record threshold for flushing file-size count data to the database.

Enhanced Ratis configuration flexibility through hdds.ratis. prefix

Ozone now allows administrators to set arbitrary Apache Ratis parameters directly within ozone-site.xml (or through Cloudera Manager safety valves) using the new hdds.ratis. prefix. When reading configuration files, Ozone automatically strips this prefix and passes the remaining key directly to the underlying Ratis server or client.

This enhancement simplifies administration and enables immediate, fine-grained performance tuning for the Ratis consensus layer such as RPC timeouts, retry policies, and log management without requiring Ozone code updates to expose new properties.

Example,

To configure a Ratis property such as raft.server.rpc.timeout.min, prepend hdds.ratis. to the property name in ozone-site.xml:

```
<property>
  <name>hdds.ratis.raft.server.rpc.timeout.min</name>
  <value>200ms</value>
</property>
```

Enhanced read scalability for Ozone Manager High Availability

Ozone Manager High Availability (OM HA) now supports serving read requests from follower nodes by leveraging Ratis Linearizable Reads. By distributing read operations across all OM nodes in the cluster rather than routing them exclusively through the leader, this feature significantly boosts read throughput for read-heavy workloads while freeing up leader capacity for write operations. Linearizable reads ensure that follower nodes maintain strict data consistency, delivering fully up-to-date metadata without stale read risks.

Increased default size for Ozone Manager Ratis Snapshots

The default value for ozone.om.ratis.snapshot.max.total.sst.size configuration property has been increased from 100 MB to 10 GB. This change provides sufficient capacity for larger datasets out of the box, improving Ozone Manager (OM) Ratis snapshot stability and performance in high-scale clusters while reducing the need for manual configuration tuning.

Periodic logging for Ozone SCM Safemode status

Storage Container Manager (SCM) now logs safemode rule status at regular intervals while in safemode. This provides continuous visibility into safemode progress and makes it easier to diagnose delayed safemode exits from logs.

The logging interval is configurable through the following configuration property:

- Property : hdds.scm.safemode.log.interval
- Default value : 1 minute

What's logged:

- **Overall state**: SCM's safemode status.
- **Prechecks**: Overall precheck completion status and count of validated precheck rules.
- **Exit rules**: Progress for each individual rule required to exit safemode.

New Grafana dashboard for Ozone SCM Safemode tracking

A new Grafana dashboard, Ozone - SCM Safemode, is now available to help you monitor the safemode status and exit progress. This is particularly useful in large clusters where safemode exit can take significant time.

The dashboard includes:

- **Rule tracking**: Individual charts for each safemode rule, showing target vs actual values to track progress toward exiting safemode.
- **Status indicator**: A visual safemode status timeline In Safemode: Displayed in red (binary value 1).
- **Exited Safemode**: Displayed in green (binary value 0).

Cloudera Runtime 7.3.2.100 CHF 1

Upgrade advisory for Ozone

Cloudera advises against installing or upgrading Ozone in Cloudera Runtime to 7.3.2.100 CHF. If you are currently running Ozone on Cloudera Runtime 7.3.2.0, Cloudera recommends skipping the 7.3.2.100 release and waiting to upgrade for an upcoming Cloudera Runtime 7.3.2.10000 (Service Pack 1).

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Ozone and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see New Features.

Cloudera Storage Optimizer

Cloudera Storage Optimizer is an intelligent data lifecycle management feature that automatically reduces storage usage by converting infrequently accessed data from replicated storage (RATIS 3×) to space-efficient erasure coding (EC). It analyzes access patterns and applies configurable policies to identify and convert cold data, reducing the storage footprint by 45% to 60% while maintaining data durability and availability. The benefit is additional usable capacity for the same licensed capacity. For more information, see Cloudera Storage Optimizer documentation.

Storage container reconciliation

By design, container replicas are identical and contain the same data. Container reconciliation is a repair mechanism that allows administrators to detect and repair container replicas that are corrupted or have otherwise diverged from each other. The data within each replica is summarized as a single data checksum. If the replica data diverges, these checksums will differ. After reconciliation, the replicas will contain the same data with matching checksums. This data checksum remains constant even as data is deleted from the container in the background. For more information, see Storage container reconciliation documentation.

Support for G1 Garbage Collector (G1GC) with JDK 17

When running Ozone on JDK 17, all Ozone server processes now default to the G1 Garbage Collector (G1GC). No action is required for most deployments.

Configuration Changes for the FIPS-compliant SASL Changes

Cloudera supports DIGEST-SHA, a new SASL mechanism, to replace DIGEST-MD5. DIGEST-SHA uses SHA256 and AES instead of MD5 and DES for message digest and encryption, respectively. As a result, DIGEST-SHA is fully FIPS-compliant. For more information, see Step 1: Prepare hosts.

Moved the port ozone.prometheus.http-port from 9094 to 9096

The port ozone.prometheus.http-port is moved from 9094 to 9096 to avoid Ozone Prometheus port colliding with HBase thrift server if they are on the same host.

Added new Ozone canaries

To configure the threshold value, the following new Ozone canaries are added along with the relevant configuration property that can be used to disable them:

- Ozone SCM heap memory canary: ozone_scm_heap_canary_enabled
- Ozone OM heap memory canary: ozone_om_heap_canary_enabled
- Ozone Datanode failed volumes canary: ozone_datanode_failed_volume_canary_enabled

Added role-specific Java options

Previously, Ozone Java options (JAVA_OPTS) were configured through a single shared setting. This change introduces separate Java option parameters for each Ozone service role, allowing more granular configuration.

The following role-specific parameters are now available:

- Storage Container Manager (SCM): ozone_scm_java_opts
- Ozone Manager (OM): ozone_om_java_opts
- Datanode (DN): ozone_dn_java_opts
- S3 Gateway (S3G): ozone_s3g_java_opts
- Recon: ozone_recon_java_opts

Each parameter applies only to its corresponding service and overrides the shared JAVA_OPTS configuration where applicable.

Added Knox proxy user configuration for Ozone Recon

You can now restrict which hosts and groups Knox can impersonate when accessing Ozone Recon by using the following configuration properties:

- ozone.recon.http.auth.proxyuser.knox.hosts: Specifies the hosts from which Knox is allowed to impersonate users .
- ozone.recon.http.auth.proxyuser.knox.groups: Specifies the user groups that Knox is allowed to impersonate.

Added new verifier for container states

The `ozone debug replicas verify` command now includes the new --container-state verifier:.

```
ozone debug replicas verify --container-state <VOLUME/BUCKET/
PATH> -o <OUTPUT.JSON>
```

This new verifier identifies keys mapped to problematic containers or replica states in SCM. It checks for the following states for every container in the replica:

- Containers: DELETED or DELETING
- Replicas: DELETED, UNHEALTHY, or INVALID

The command generates a JSON file containing details on keys, blocks, and replicas, along with a pass or fail status for each check. It supports the following options:

- --all-results: Displays all verification outcomes
- -o: Specifies the output file path
- --container-cache-size: Sets the number of containers stored in the in-memory cache for the --co ntainer-state check.

  The default value is 1 million containers (approximately 43MB of memory). The value must be greater than zero, otherwise, the system uses the 1 million default.

  Important: The system ignores this option if you do not use the --container-state flag.

Enhanced volume health checks triggered by container scanner

When a container scanner (background or on-demand, data or metadata) marks a container as unhealthy due to corruption, the system now automatically triggers an on-demand scan of the associated volume. This enhancement helps detect and address broader volume-level issues that might exist beyond the affected container, improving overall data integrity and reliability.

New container health metrics

In Cloudera Base on premises 7.3.2.0 and higher releases, Ozone Recon now exposes additional metrics for the Container Health Task to improve observability of container state across the cluster. Recon tracks the number of containers in the following states:

- Missing
- Under-replicated
- Over-replicated
- Mis-replicated

Increased default Ratis segment size for Ozone Manager (OM) and Storage Container Manager (SCM)

To improve startup performance, the default values of the ozone.om.ratis.segment.size and ozone.scm.ha.ratis.segment.size configurations are increased from 4 MB to 64 MB.

- The ozone.om.ratis.segment.size configuration sets the size of the raft segment used by Apache Ratis on OM.
- The ozone.scm.ha.ratis.segment.size configuration sets the size of the raft segment used by Apache Ratis on SCM.

Renaming all legacy HDFS configuration keys and the corresponding Java constants to HDDS

All legacy configuration keys previously prefixed with dfs. (used by HDFS) are updated to use the hdds. prefix, aligning them with the HDDS component. Backward compatibility is maintained. The system will continue to recognize the old dfs. keys, but they are now deprecated. Deprecation notices and mappings are handled internally by adding DeprecationDelta in OzoneConfiguration.

Additionally, the corresponding Java constants for both the configuration key and the default value are renamed accordingly.

All deletion configurations are dynamically configurable without restart

The following deletion configurations are dynamically reconfigurable without requiring a restart:

- For Ozone Manager: ozone.thread.number.dir.deletion
- For Datanode:
  - ozone.block.deleting.service.interval
  - ozone.block.deleting.service.timeout

Support for callback on completed reconfiguration

The reconfiguration framework now supports a reconfiguration completed event. Components can now register callbacks that trigger after all properties in a task are processed. This allows the

system to validate and apply interdependent settings automatically. You can access the status and values of the reconfigured properties using the ReconfigurableBase#getReconfigurationTaskStatus method within the callback.

Moved and enhanced Ratis Log Parsing command in Ozone

The `ozone debug ratislogparser` command is moved under the new `ozone debug ratis parse` subcommand. The parse subcommand now supports parsing different Ratis files, providing a more general and extensible interface.

The new --role=[om, scm, datanode] optional flag is added to specify which schema to use when parsing logs. By default, the schema is set to generic.

Enhanced output for `ozone debug replicas chunk-info` command

The `ozone debug replicas chunk-info` command now provides detailed information about the chunks associated with a given key. Previously, chunk-related details were only available when the --verbose flag was used. With this update, the command includes the BlockData information retrieved from the getBlock call, along with additional details about each replica. The JSON output structure is also improved for clarity.

Unified JSON output for `ozone debug replicas verify` command

The `ozone debug replicas verify` command now outputs JSON information for each key and the checks performed on it. As new verification checks are added, their results are included in the same JSON objects. By default, the output is skipped for keys that pass all specified checks, reducing unnecessary output.

Increased Ozone Manager (OM) default snapshot limit

The default value for the ozone.om.fs.snapshot.max.limit Ozone Manager property is increased, allowing support for up to 10,000 snapshots by default (previously 1,000). This change enables you to create and manage a significantly larger number of snapshots without requiring manual configuration changes.

Ozone Recon UI: Added new Cluster Health panel to the Grafana Dashboard

A new Cluster Health panel is added to the existing Ozone - Overall Metrics Grafana dashboard. This panel displays the cluster growth rate, calculated dynamically using Prometheus metrics. By default, the panel displays the growth rate over the last 1 hour, but you can configure the time window as required.

Renamed Ozone configuration property ozone.compaction.service.enabled

The ozone.compaction.service.enabled configuration property is renamed to ozone.om.compaction.service.enabled. This change aligns with the Ozone configuration key conventions, as the property is used only by the Ozone Manager (OM).

Added bucket layout flag to `freon rk` command

Previously, the `freon randomkeys (rk)` command did not allow you to specify the bucket layout, always creating buckets with the default layout. With this improvement, a new option is introduced to set the bucket layout using the command, providing greater flexibility.

Enhanced DataNode listing with usage-based sorting

The `ozone admin datanode list` command now supports the new --most-used and --least-used CLI options. These options enable you to sort Datanodes based on their usage. When used, the command output includes additional details such as used space, total capacity, and percentage of space used for each Datanode. This enhancement helps you identify the most and least used Datanodes in your Ozone cluster.

Updated the Ozone SCM container placement policy description

The descriptions for the following Ozone configuration properties are corrected in the ozone-default.xml file to clarify that the default placement policy is SCMContainerPlacementRackAware, instead of SCMContainerPlacementRandom:

- ozone.scm.container.placement.impl

- ozone.scm.container.placement.ec.impl

Additionally, if the ozone.scm.container.placement.impl property is left unset, the code fallback logic now uses the SCMContainerPlacementRackAware policy to match the default value in the ozone-default.xml file.

Enhanced container creation command with replication options

The `ozone admin container create` command now supports additional replication configurations. Previously, containers were limited to a replication factor of ONE and a replication type of STANDALONE. Administrators can now specify replication types (RATIS and EC), replication factors, and algorithms when creating containers using the following parameters:

```
ozone admin container create --replication=<REPLICATION> --repli
cation-type=<TYPE>
```

What's New in Apache Phoenix

New features and functional updates for Phoenix are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

Phoenix supports TLS 1.3

Phoenix now supports Transport Layer Security (TLS) version 1.3 for encrypting communication between clients and the Phoenix Query Server. TLS 1.3 provides improved security and performance over TLS 1.2 by using modern cipher suites such as TLS_AES_128_GCM_SHA256 and TLS_AES_256_GCM_SHA384. You can configure the TLS version and cipher suites for Phoenix Query Server connections through Cloudera Manager.

For more information, see Configuring the TLS versions for Phoenix.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces new features of Phoenix and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all updates in Cloudera Runtime 7.3.1.x, see New Features.

IPv6 support for Phoenix

Starting with the 7.3.2 release, Phoenix client supports IPv6 with dual-stack functionality, allowing seamless communication over both IPv4 and IPv6 networks. This capability improves network scalability, future-proofs deployments, and enhances overall platform security.

For more information, see Enabling IPv6 support for Phoenix.

Phoenix supports cluster-wide metadata upgrade blocks

Phoenix now provides the ability to block cluster-wide metadata auto-upgrades. To prevent an upgrade from occurring, you must add a BLOCK_UPGRADE row in the SYSTEM.MUTEX or SYSTEM:MUTEX HBase table (depending on whether namespace mapping is enabled), in the

0:MUTEX_VALUE column family, before you start Phoenix with the upgraded versions. To add the row by using the HBase shell, run the following command:

```
put 'SYSTEM:MUTEX','BLOCK_UPGRADE','0:MUTEX_VALUE','MUTEX_LOCKED'
```

For a Java code example, see the LoadSystemTableSnapshotBase.java file in the Apache Phoenix repository.

Phoenix JDK 17 upgrade

Phoenix only supports JDK 17 starting with the 7.3.2 release; upgrading to JDK 17 ensures your application stays competitive and maintainable by providing improved performance, increased security, modern language features, and long-term support.

Deprecation of OMID service

The OMID service was deprecated in Cloudera Runtime 7.3.2 and will be removed in a future release. If you still require OMID, you can install the service through Cloudera Manager.

Apache Phoenix component upgraded to 5.2.1

The Phoenix runtime component upgrades from 5.1.3 to 5.2.1. This release delivers performance enhancements and introduces cluster-wide metadata upgrade blocks, providing you with greater control during upgrades and increasing stability in modern environments.

What's New in Apache Ranger

Learn about the new features of Apache Ranger in Cloudera Runtime 7.3.2.0, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

New enhanced policy creation interface in Ranger

Apache Ranger now provides an enhanced policy creation interface that enables you to create policies across available services in a streamlined, single workflow. The new interface uses a step-by-step guided flow, reducing the manual effort of navigating to each individual service's policy listing page to create a policy. You can navigate between steps non-linearly, allowing you to jump to any step in the workflow to review or modify your selections. The interface also supports editing existing policies using the same guided experience. For more information, see

For more information, see Creating any policy from the policy creation form.

CDP Observability integration with Ranger Admin

Ranger Admin now supports integration with CDP Observability using OpenTelemetry. The cdp-opentelemetry-collector pulls metrics exposed by Ranger Admin at a 1-minute interval, enabling monitoring and observability for Ranger. This integration is controlled by the ranger.admin.obs ervability.metrics.enabled property in ranger-admin-site.xml and is enabled by default. To disable metrics collection, set the property to false via Cloudera Manager # Ranger # Configuration # Ranger Admin Advanced Configuration Snippet (Safety Valve) for conf/ranger-admin-site.xml.

For more information, see Overview of Ranger and Observability integration.

TLS 1.3 support

Ranger supports TLS 1.3 for stronger security and compliance.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features in this release.

Cloudera Runtime 7.3.2.0

The new features for Apache Ranger in Cloudera Runtime 7.3.2.0 include cumulative updates from previous releases. This version specifically incorporates features introduced in Cloudera Runtime 7.3.1.100 through 7.3.1.706. For a complete list, see New Features.

Support for Amazon S3-compatible object stores through RAZ

Support has been added for Amazon S3-compatible object stores through the Ranger Remote Authorization Service (RAZ).

Cloudera Base on premises defaults to HDFS and natively supports the Ozone object store. There was limited support for Amazon S3-compatible object stores with respect to executing workloads like Hive and Spark. Not all workloads were supported from an authorization perspective on Amazon S3-compatible object stores. The inconsistent policy framework created challenges for the Policy Administrator in terms of applying corporate policies to secure sensitive data. Managing data access across teams and individuals was an architectural challenge.

The Ranger RAZ resolves this challenge by using Apache Ranger policies to authorize access to Amazon S3-compatible object storage, similar to HDFS files. For details, see Access Control for Amazon S3-compatible object stores.

Ranger has been upgraded

Ranger has been upgraded from version 2.4.0 to 2.6.0, incorporating the latest fixes and improvements from the Ranger project while maintaining compatibility with existing Cloudera deployments.

Ranger HBase plugin optimizations

You can now enable column authorization optimization, wherein column authorization is skipped when a user is fully authorized for the column families. To enable this optimization, add the following Ranger service configuration for cm_hbase from Ranger UI: ranger.plugin.hbase.colu mn.auth.optimized=true. The new configuration can lead to significantly better performance of multiget and multiput workloads, wherein thousands of columns may be accessed together. No HBase service restart is required to enable or disable this configuration.

When the optimization is enabled, there is a behavior change in auditing in Ranger. For more information, see Behavioral Changes in Apache Ranger.

Ranger RMS supports for multiple storage types

Starting with Cloudera Base on premises 7.3.2.0 release, RMS supports S3 storage in addition to HDFS and Ozone, allowing it to track mapping data for all these storage types simultaneously.

After the first full-synchronization run, RMS downloads mappings for tables and databases present in the Hive Metastore (HMS). These tables/databases could be located in S3, OZONE, and HDFS file systems. RMS will map the tables and databases with their respective storage locations and store them in the Ranger database with their associated service ID (cm_s3, cm_ozone, or cm_hdfs). When raz-s3-chained-plugin, ranger-ozone-plugin, or ranger-hdfs-plugin requests mappings, it only retrieves the mappings relevant to its service. Specifically, raz-s3-chained-plugin downloads mappings exclusively for tables/databases whose storage location is S3, ranger-ozone-plugin for Ozone, and ranger-hdfs-plugin for HDFS.



Important: For users who had RMS configured for HDFS/Ozone before Cloudera Base on premises 7.3.2.0 and also have external table data stored in S3, a full synchronization is necessary after upgrading. This ensures that any databases and tables created prior to the upgrade, and whose data resides in S3, are correctly mapped. Any newly created tables or database mappings will be synchronized in RMS.

User's last login timestamp visibility

The Ranger Admin Web UI now displays the logged-in user's last login date and time on all pages.

Earlier, this information was only available under Ranger Audits Login Sessions .

Now the user's last login timestamp is persistently visible across all pages. For a user's initial login, the last login timestamp will not be displayed. Upon subsequent logins, the timestamp of the preceding successful login will be displayed.

Storage change for Ranger admin audit logs

Starting from Cloudera Runtime 7.3.2.0, Ranger admin audit logs are stored in the x_trx_log_v2 table in the Ranger database. In releases earlier than 7.3.2.0, Ranger admin audit logs were stored in the x_trx_log table in the Ranger database. Hence, if you upgrade to Cloudera Runtime 7.3.2.0 from any previous version and you want to retain your old data, you must migrate the Ranger admin audit logs from the x_trx_log table to the x_trx_log_v2 table.



Note: In the case of a fresh installation of Cloudera Runtime 7.3.2.0, this migration process is not required.

For more information on how to migrate the data, see Migrating Ranger admin audit logs.

Back up ranger_audits Solr collection

When upgrading to Cloudera Runtime 7.3.2.0 and later releases that update the Ranger Audits Solr schema, Cloudera Manager runs an internal service command to apply the schema changes safely.

Although the command is designed to be non-destructive, Cloudera strongly recommends creating a backup of the Ranger Audits Solr collection before starting the upgrade. This ensures a quick restoration of audit data if anything goes wrong during the upgrade.

For more information, see Back up ranger_audits Solr collection.

Ranger mixed case group comparison

When Ranger Usersync is configured with case conversion and special character replacement using Regular Expression (regex), Ranger Usersync transforms the original user or group names from the source, for example, AD or LDAP, before storing them in the Ranger Admin database. Previously, if a Ranger plugin used the original name during authorization checks, the check failed because the Ranger Admin only recognized the transformed name.

This issue is now fixed. The fix is configurable at the plugin level using the ranger.plugin.<servi ceType>.supports.name.transformation property, allowing users to enable or disable transformation based on their environment needs. For more information, see Handling inconsistent username and group name conventions for consistent authorization.

ranger.plugin.hive.urlauth.filesystem.schemes configuration with default value of hdfs:

The default value of the ranger.plugin.hive.urlauth.filesystem.schemes configuration is hdfs:,fi le:,wasb:,adl:. When the ranger.plugin.hive.urlauth.filesystem.schemes configuration includes hdfs: as a default filesystem scheme in both HMS and HiveServer2, and you perform a CREATE EXTERNAL TABLE operation with a LOCATION clause, if both the Ranger Hive and HDFS plugins are enabled, the Ranger Plugin Authorizer performs authorization checks on all files within the target directory. This results in performance degradation that scales up linearly with the number of files.

The default value of the parameter remains unchanged, but a warning message has been added for the user. For more information, see Hive authorizer URL policy with hdfs default value.

Added validation for whitespace in Ranger policy names

Ranger Admin now validates and handles leading and trailing whitespace in policy names and related properties (for example, db_name vs db_name) created through the Ranger API. Policies that contain unintended whitespace are detected and normalized so they no longer appear inconsistent between the API and the Ranger Admin web UI, helping prevent confusing behavior and reducing issues for application teams.

Bulk delete policies using policy name prefix

Added a Ranger REST API endpoint that supports bulk deletion of policies by policy name prefix (wildcard). Administrators can now delete multiple policies in a single request by specifying a

service name and a policy name prefix (for example, JAMES*), improving operational efficiency over deleting policies one by one.

Support SASL bind for Ranger Usersync with AD/LDAP

Introduced support for SASL GSSAPI authentication in Ranger Usersync when connecting to AD/LDAP. This allows customers to avoid using simple bind credentials and mitigates performance and reliability issues observed with SSSD on RHEL 8.8 (for example, incomplete getent results with enumerate=true).

Ranger RAZ service-admin delegation support

Ranger RAZ now supports delegation for users with service-admin privileges. When making authorization calls to RAZ, users defined in the Ranger service configuration as service.admin.users or belonging to groups in service.admin.groupscan authorize access on behalf of other users. Requests from these service-admin users are also trusted for details such as the resourceOwner, improving support for administrative and service-level access workflows.

RAZ to support authorization for CAII service on Cloudera Base on premises

In Cloudera Base on premises 7.3.2.0, Ranger RAZ now supports the Cloudera AI Inference (CAII) service for authorization. During initialization, Ranger automatically creates the CAII service definition and a corresponding service instance, similar to other RAZ#managed services (for example, cm_hive and cm_s3). RAZ has been enhanced to authorize access to CAII resources and to honor service.admin.users and service.admin.groups so that designated service administrators can authorize access on behalf of other users.

Improved handling of Kafka super users in Ranger Kafka plugin

The Ranger Kafka authorizer now automatically parses the Kafka super.users configuration property and adds those users to the plugin's super-user list. This allows Strimzi and other deployments that rely on super.users (for example, cluster operator and broker users) to have the expected unrestricted access to Kafka resources without requiring duplicate super-user configuration in Ranger.

Ranger Usersync LDAP referral handling improved

In Cloudera Runtime 7.3.2.0, the default value of the Ranger Usersync LDAP referral configuration has been changed from ignore to follow (ranger.usersync.ldap.referral=follow). This enables Ranger Usersync to follow LDAP referrals during group searches, which helps prevent javax.naming.PartialResultException: Unprocessed Continuation Reference(s) errors commonly seen with Active Directory LDAP and reduces the need for relying on the AD Global Catalog for referral resolution.

IP-based access control for policies

Ranger now supports IP-based access control across additional services by using context enrichers and policy conditions. Administrators can configure IP location data files, register the RangerFileBasedGeolocationProvider context enricher, and define IP-based policy conditions (for example, accessed-from-ip) to allow or deny access based on client IP ranges. This capability is not enabled by default and must be explicitly configured, following the steps described in the Apache Ranger Geolocation based policies documentation.

Improved REST API responses for policy?policyNamePartial

Ranger policy REST API responses for policy?policyNamePartial have been enhanced to optionally include policy metadata fields such as createTime and updateTime. These meta attributes are now configurable for retrieval, allowing administrators to access policy creation and modification timestamps directly from the public v2 policy API when enabled.

Ranger Swagger API usability improvement

The Ranger /service/xusers/users Swagger API now supports the optional query parameters syncSource and userRole. These parameters are exposed as fillable fields in Swagger to make it easier to construct complex user search requests (for example, filtering users by sync source such as Unix or LDAP/AD, or by role such as admin), especially for less technical users and environments with large user populations.

Improved authzmigrator script

Improved the authz-export.sh (authzmigrator) script so it no longer relies on a hard-coded Cloudera parcel path and JDBC driver. The script now uses the configured Cloudera Manager parcel directory (via SENTRY_HOME) and automatically selects the appropriate JDBC driver for the underlying database, preventing NoClassDefFoundError: javax/jdo/JDOHelper and making it more robust in environments with custom parcel layouts and non-default DB flavors.

Downlist the public group from Select Group option in New Policy page

To reduce the risk of accidentally assigning the Public group when creating a new Ranger policy, the Tab key no longer selects a group in the Select Group field on the New Policy page. Pressing Tab now only moves focus to the next field as expected, while the Enter key remains the standard way to confirm a group selection. This change significantly lowers the chance of unintended Public group assignment without impacting normal group selection behavior.

Delete Ranger RAZ access logs based on max number of files

To prevent Ranger RAZ access logs from filling up disk space on busy environments, Ranger RAZ now supports log rotation based on a configurable maximum number of access log files. Previously, access logs could only be rotated based on age (with a default of 15 days), which could still allow logs to exhaust the volume on high-activity clusters. With this enhancement, administrators can limit the number of retained Ranger Raz access log files, helping ensure Cloudera Manager and other services remain available even under heavy load.

Ranger access audit option to exclude resource

Improved Ranger access audit filtering to support excluding specific resources by name.You can now use an Exclude Resource Name filter in the Ranger Access audit UI so that audit searches can omit events for specified resources, leveraging existing backend Solr support for these exclusions.

Included authz_export.tar.gz for Sentry to Ranger migration with Cloudera Manager package

For CDH#to#CDP upgrades that involve Sentry#to#Ranger migration, the authz_export.tar.gz package is now bundled with the Cloudera Manager installation, in the same location as authz-in gest.zip. This eliminates the need to contact Cloudera support to obtain the export archive from internal sources and simplifies authorization policy migration.

Ranger Admin logs to show the local timezone rather than UTC

Ranger Admin now records log timestamps in the local system timezone of the host instead of UTC. This aligns Ranger Admin logs with other Ranger components (such as UserSync and TagSync) and removes the need for manual log4j configuration, making troubleshooting and correlation with other system logs easier.

Cloudera Manager now sanitizes cluster names

Cloudera Manager now sanitizes cluster names that contain spaces when creating and evaluating Ranger plugin repositories that use the UNIQUE_PER_SERVICE naming strategy. Spaces in the derived Ranger repository (service) name are automatically replaced with underscores (_), ensuring compatibility with the Ranger constraint introduced in RANGER-2808 that disallows spaces in service names.

Improved Ranger Admin Diagnostic collection command from Cloudera Manager scripts

In Cloudera Manager 7.13.2.0, the Ranger Admin diagnostic collection command has been enhanced. A new configuration option, ranger.admin.diag.metrics.collection.type, allows you to control which types of Ranger metrics are collected during the Cloudera Manager diagnostics command. This helps avoid long#running Ranger Admin metrics collection (for example, cases where collection could exceed 40 minutes and cause the Cloudera Manager diagnostic command to stop) by letting you limit or skip specific metrics types as needed.

Updated Ranger RMS, TAGSYNC, and USERSYNC CSDs

Updated Ranger RMS, TAGSYNC, and USERSYNC Cloudera Service Descriptors (CSDs) to support the new Ranger RMS high availability (HA) implementation used by the ranger-commo n-ha module. The RMS server and tagsync/usersync service port configs are now managed as

first#class CSD properties rather than via safety valves, improving HA configuration consistency and upgrade handling.

Cloudera Manager updates Ranger plugin service repositories with the correct cluster configurations

Cloudera Manager now updates Ranger plugin service repositories with the correct cluster configurations for test connection and resource lookup. For Ranger plugins on HBase, Knox, Atlas, and Ozone, Cloudera Manager automatically calculates and populates the required URLs, principals, and related settings instead of creating repositories with default, nonfunctional values. This reduces configuration errors and addresses frequent “Test connection” and lookup failures seen in customer deployments.

Enforce truststore configuration when Ranger plugin is enabled with Ranger admin SSL

When the Ranger plugin is configured to use Ranger Admin over SSL, Cloudera Manager now enforces truststore configuration for the plugin services. If the truststore settings are missing, the Ranger Plugin Services page displays a validation warning, helping prevent plugin sync failures caused by an unconfigured truststore.

Improved control of Ranger audit spooling to prevent disk exhaustion

Added Cloudera Manager configuration options for the Ranger plugins to control audit file spooling, including the spool file rollover interval and the maximum number of archived spool files. These settings are now exposed centrally for HDFS and Solr Ranger audits, reducing the risk of /var/log filling up when the audit destination is unavailable and eliminating the need for per-service safety#valve configuration.

Ranger audits collection creation reliability improved

Updated the Ranger audit configuration so that the ranger.audit.solr.time.interval setting now uses a 60-second (60000 ms) delay between retries when creating the ranger_audits collection in Solr, improving reliability in environments where Solr starts later than Ranger.

What's New in Ranger KMS

Learn about the new features of Ranger KMS in Cloudera Runtime 7.3.2.0, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

Ranger KMS supports new Key Derivation Functions (KDFs) and ciphers

Ranger KMS now supports configurable Key Derivation Functions (KDFs) and ciphers for Master key and Zone key encryption. KMS uses AES-256 for all key generation and allows administrators to configure cryptographic parameters for enhanced security. To understand and configure the supported KDFs and ciphers Configuration of cryptographic parameters.

TLS 1.3 support

Ranger KMS supports TLS 1.3 for stronger security and compliance.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features in this release.

Cloudera Runtime 7.3.2

Ranger KMS in a federated deployment

Ranger KMS can now be deployed in a federated cluster for key management.

A data cluster is a cluster where application and data processing occur. The cluster stores and processes actual datasets but does not directly manage encryption keys. A security cluster (with the Ranger KMS service installed) is managed separately from the data cluster. It manages the key lifecycle operations (for example, generation, rotation, storage). This separation of tasks enhances security by isolating the administration of the data and security clusters.

For more details, see Installing Ranger KMS in a federated deployment.

What's New in Schema Registry

New features and functional updates for Schema Registry are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new features in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

There are no new features introduced in this release.

What's New in Apache Solr

New features and functional updates for Solr are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

Solr supports incremental backup and restore

Solr now supports incremental backup and restore through the Collections REST API. Incremental backups optimize storage and reduce data transfer by backing up only the changes since the previous backup. You can create incremental backups, restore specific backup versions, and list available backups by using the Collections REST API `BACKUP`, `RESTORE`, and `LISTBACKUP` commands. Incremental backup is the default behavior for API-based backups. To create a traditional non-incremental backup through the API, set the incremental=false parameter on your backup request. Backups created through snapshots by using the solrctl utility remain non-incremental.

TLS 1.3 support for Solr

Cloudera Runtime 7.3.2.10000 introduces TLS 1.3 support for Solr. Two new configuration parameters are available in Cloudera Manager that you can use to control the TLS protocol versions and cipher suites for Solr SSL/TLS communication:

- Solr TLS/SSL Enabled Protocol Versions (solr.jetty.ssl.enabledProtocols / solr_ssl_enabled_protocols) - Specifies the TLS protocol versions that Solr accepts for SSL/TLS connections. You can select TLSv1.2, TLSv1.3, or both TLSv1.2 and TLSv1.3.

- Solr TLS/SSL Cipher Suites (solr.jetty.ssl.cipherSuites / solr_ssl_cipher_suites) - Specifies the cipher suites that Solr uses for TLS-encrypted connections. When you enable TLS 1.3, Cloudera recommends that you configure cipher suites that are compatible with the selected protocol version.

For more information, see Configure TLS/SSL encryption for Solr.

TLS 1.3 support for Key-Value Store Indexer

Cloudera Runtime 7.3.2.10000 introduces TLS 1.3 support for the Key-Value Store Indexer (KS_INDEXER) in Solr. Three new configuration parameters are available in Cloudera Manager that you can use to secure RPC communication between the HBase Indexer and HBase:

- Enable TLS/SSL for HBase-Indexer TLS enable (hbase.server.netty.tls.enabled / hbase_indexer_tls_server_enabled) — Enables TLS/SSL for the HBase Indexer server using the Netty transport layer. Because the indexer operates similarly to an HBase RegionServer, this parameter behaves the same as the corresponding HBase TLS parameter.
- Supported SSL/TLS versions (hbase.rpc.tls.enabledProtocols / hbase_indexer_rpc_tls_versions) — Specifies the TLS protocol versions that the HBase Indexer accepts for RPC connections. You can select TLSv1.2, TLSv1.3, both TLSv1.2 and TLSv1.3, or copy the value from the global settings.
- TLS Cipher List (hbase.rpc.tls.ciphersuites / hbase_indexer_rpc_tls_ciphers) — Specifies the cipher suites that the HBase Indexer uses for TLS-encrypted RPC connections.

For more information, see Configure Lily HBase Indexer to use TLS/SSL.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features in this release.

Cloudera Runtime 7.3.2:

Solr JDK 17 upgrade

Solr supports JDK 17 starting with the 7.3.2 release. Upgrade to JDK 17 to gain improved performance, increased security, modern language features, and long-term support, ensuring your application remains competitive and maintainable.

What's New in Spark

Learn about the new features of Spark in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

Rebase on Apache Spark 3.5.8

Spark shipped with Cloudera Runtime 7.3.2.10000 SP1 is based on Apache Spark 3.5.8 (previously 3.5.4 in Cloudera Runtime 7.3.2). This maintenance rebase includes bug fixes and security updates from upstream Apache Spark. For more information, see the Apache Spark documentation.

Important:

When Cloudera rebases the Apache Spark version shipped with or CDS, minor and maintenance releases can still change SQL behavior, configuration defaults, and serialization. Recompile Java and Scala applications built against Spark, and retest PySpark and Spark SQL workloads in a non-production environment before you upgrade production clusters.

Review *Upgrading Spark* and *Migrating Spark applications* for steps that match your source and target Spark versions.

Read the *Apache Spark migration guides* for the **exact** Apache Spark version on your cluster (for example 3.5.8 or 3.5.4).

See Upgrading Apache Spark 3.5.4 to Spark 3.5.8 and Migrating Spark applications.

FIPS 140-3 support for Spark 3

Spark 3 supports FIPS 140-3 on FIPS-enabled clusters in this service pack. Spark cryptographic modules and Cloudera SASL are updated for FIPS 140-3 validation.

For configuration details, see Spark and Livy on FIPS 140-3 clusters.

Encryption in transit for Spark

Spark 3 supports stronger encryption in transit when you enable Encrypt all ports in Cloudera Manager, including Spark network encryption, HTTPS-only History Server and UI listeners, and updated IO encryption defaults. Spark History Server and UI TLS can use TLS 1.3 when your Cloudera Manager TLS policy allows it.

For configuration and behavioral changes, see Encryption in transit for Spark and Livy in Cloudera Runtime 7.3.2.10000 SP1, Configuring TLS/SSL encryption manually for Spark, and Behavioral Changes in Spark.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2.0

Rebase Spark3 to Apache Spark 3.5.4 in Cloudera Runtime.

Spark 3.5.4 is the default Spark version in Cloudera Runtime. Refer to the Apache Spark documentation for more information on changes.

Important: Refer to Migrating Spark applications for more information on migrating your existing Spark applications.

Spark 3 is the default in Cloudera Runtime

Spark 3 is the default Spark version in Cloudera Runtime. Spark 2 has been removed and no longer available in 7.3.2.0

Important:

Spark 3 contains a large number of changes from Spark 2.

Refer to Upgrading Spark for more information on upgrading Spark clusters to 7.3.2.0, and Migrating Spark applications for more information on migrating your existing Spark applications between versions 2 and 3.

Related Information

Upgrading Spark

Migrating Spark applications

What's New in Spark Atlas Connector

New features and functional updates for Spark Atlas Connector are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no Spark Atlas Connector-specific features in this service pack. Lineage metadata is sent to Apache Atlas and Kafka using the TLS and Kerberos settings for those services. On FIPS 140-3 clusters, the connector does not ship a separate FIPS cryptographic module.

For Spark and Livy encryption in transit, see .

For Apache Atlas TLS 1.3, see Configuring TLS/SSL for Apache Atlas and What's New in Apache Atlas.

For FIPS configuration, see .

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

There are no new features introduced in this release.

What's New in Sqoop

New features and functional updates for Sqoop are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

There are no new features introduced in this release.

What's new in Streams Messaging Manager

New features and functional updates for Streams Messaging Manager are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

TLS metrics delivery to Cloudera Service Monitor

Streams Messaging Manager sends metrics to Cloudera Service Monitor over TLS when TLS is enabled for Cloudera Management Service. Restart Streams Messaging Manager after you enable or disable TLS for Cloudera Service Monitor. For information about enabling TLS on Cloudera Management Service roles, see *Configuring TLS/SSL for Cloudera Management Service*.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

There are no new features introduced in this release.

What's New in Streams Replication Manager

New features and functional updates for Streams Replication Manager are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new features in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

Reverse Checkpointing

Streams Replication Manager now supports reverse checkpointing. This feature enables the tracking and replication of consumer offsets from a target cluster back to a source cluster. By tracking offsets in the reverse direction, you ensure that the progress made by consumer groups on a backup cluster is preserved and translated back to the primary cluster during a failback scenario.

Reverse checkpointing minimizes message duplication upon failback by mapping the offsets from the replica topic back to the equivalent offsets in the source topic. To enable this feature, you must configure the following in Cloudera Manager:

- Set the cloudera.reverse.checkpointing.enabled property to true.
- Enable bidirectional replication in the Streams Replication Manager's Replication Configs property.

In addition to service configurations, you must use the srm-control tool to explicitly allowlist topics for reverse checkpointing using the reverse-checkpointed-topics command. Consumer group replication must also be enabled in both directions.



Important: Reverse checkpointing relies on cluster prefixes to identify replica topics and is therefore not compatible with the IdentityReplicationPolicy.

Single REST server for all replication flows

Streams Replication Manager now uses a single REST server with a single port to handle inter-worker communication for all replication flows. Previously, a dedicated REST server was started for each replication flow. The new implementation exposes only the endpoints required for inter-worker coordination and task configuration updates. These endpoints are restricted to inter-worker communication and cannot be accessed externally. The legacy per-flow REST server implementation is deprecated in 7.3.2 and will be removed in a future release. Cloudera recommends that you migrate your Streams Replication Manager clusters to the new implementation.

Suppressing internal metrics topics

You can now configure the Streams Replication Manager Service to suppress the eager creation of srm-metrics topics for all possible replication flows. This prevents the creation of unused topics. To enable this behavior, set the metrics.topic.creation.for.possible.flows.enabled property to false.

Configurable timeout for Streams Application Kafka Connection Health Test

A new SRM Service Streams Application Connection Test Timeout (streams.replication.manager.service.streams.application.connection.test.timeout) Cloudera Manager configuration option is now available for the Streams Replication Manager Service. It sets the timeout, in seconds, for the Streams Application Kafka Connection Health Test, which periodically checks connectivity to the target Kafka cluster. The default is 1 second.

What's new in YARN and YARN Queue Manager

New features and functional updates for YARN and YARN Queue Manager are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

The following new properties and configyrations are introduced in this release:

MapReduce security allowed users property

Defines a list of trusted users permitted to bypass the denied-tasks list. Designed for system-level or administrative workflows requiring restricted task implementations, this setting skips blocklist validation for designated submitters while maintaining standard Hadoop ACL and authentication checks.

Table 7: Property name and its details

Property name	Property type	Default	Example
mapreduce.security.allowed-users	List of usernames	Empty	alice, bob

MapReduce task-level security enforcement

When enabled, the Application Master performs validation of user-submitted mapper, reducer, and other task-related classes before launching containers. This mechanism protects the cluster from running disallowed or unsafe task implementations as defined by administrator-controlled policies.

Table 8: Property name and its details

Property name	Property type	Default
mapreduce.security.enabled	boolean	false (security disabled)

MapReduce Security Property Domain

The property domain specifies the MapReduce configuration keys for user-defined class names, including mappers, reducers, and partitioners, that are executed during task runs. The Application Master inspects these property values to verify that no specified class matches the list of denied tasks. System administrators can customize this list to broaden or narrow the validation scope.

Table 9: Property name and its details

Property name	Property type	Default
mapreduce.security.property-domain	list of configuration keys	map.sort.classL

List of new configurations:

- mapreduce.job.classloader.system.classes
- mapreduce.job.combine.class
- mapreduce.job.combiner.group.comparator.class
- mapreduce.job.end-notification.custom-notifier-class
- mapreduce.job.inputformat.class
- mapreduce.job.map.class
- mapreduce.job.map.output.collector.class
- mapreduce.job.output.group.comparator.class
- mapreduce.job.output.key.class
- mapreduce.job.output.key.comparator.class
- mapreduce.job.output.value.class
- mapreduce.job.outputformat.class
- mapreduce.job.partitioner.class
- mapreduce.job.reduce.class
- mapreduce.map.output.key.class
- mapreduce.map.output.value.class

TLS 1.3 support

Ranger supports TLS 1.3 for stronger security and compliance.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

Hadoop rebase summary

In Cloudera Runtime 7.3.2, Apache Hadoop is rebased to version 3.4.1. The Apache Hadoop upgrade improves overall performance and includes all the new features, improvements, and bug fixes from versions 3.2, 3.3, and 3.4.

Table 10: Improvements added from Apache Hadoop 3.2 to 3.4 versions

Apache Hadoop version	Apache Jira	Name	Description
3.4	YARN-9279	YARN Hamlet Package Removal	The deprecated org.apache.hadoop.yarn.webapp.hamlet package is now completely removed to improve maintainability. This is an incompatible change in Hadoop YARN 3.4.0+. Applications relying on this old package must be updated to use the org.apache.hadoop.yarn.webapp.hamlet2 package. This affects the YARN webapp component.
3.4	YARN-10820	Enhanced Reliability for YARN node list Command	The thread-safety issue is fixed in GetClusterNodesRequestPBImpl, that previously caused intermittent failures, such as java.lang.ArrayIndexOutOfBoundsException, with the YARN node list command. This change affects the YARN client in Hadoop YARN 3.4.0, 3.3.2, and 3.2.4, thereby, eliminating random crashes when running the YARN node list command.

Table 11: Issues fixed between Apache Hadoop versions 3.2 to 3.4

Apache Hadoop version	Apache Jira	Name	Description
3.3	MAPREDUCE-6990	MapReduce task initialization Timeout issue	Previously, MapReduce jobs stopped responding if a task terminated before sending its first heartbeat, as the task never timed out and remained stuck indefinitely in a "STARTING" state. This issue is now resolved by introducing a dedicated timeout mechanism specifically designed to catch and terminate tasks that fail to initialize and send their first heartbeat.
3.4	YARN-9809	Miscommunication between RM and NM when NodeManagers are unhealthy	Previously, if a NodeManager (NM) was registered in an unhealthy state, it did not communicate the status immediately. As a result, the Resource Manager (RM) mistakenly scheduled many containers to that unhealthy node before the first heartbeat was received. Once the first heartbeat finally arrived, the RM recognize the unhealthy status and abruptly ended all the recently scheduled containers, causing unnecessary task failures and wasted resources. This issue is now resolved and NMs now explicitly supply their health status during their initial registration with the RM.

What's New in Apache ZooKeeper

New features and functional updates for ZooKeeper are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new features in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new features in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new features introduced in this release.

Cloudera Runtime 7.3.2

IPv6 support for ZooKeeper

Starting with the 7.3.2 release, ZooKeeper supports IPv6 with dual-stack functionality, allowing seamless communication over both IPv4 and IPv6 networks. This capability improves network scalability, future-proofs deployments, and enhances overall platform security.

ZooKeeper JDK 17 upgrade

ZooKeeper supports JDK 17 starting with the 7.3.2 release. Upgrade to JDK 17 to gain improved performance, increased security, modern language features, and long-term support, ensuring your application remains competitive and maintainable.

IPv6 prefixes support for ZooKeeper ACLs

Support for IPv6 prefixes is now added to ZooKeeper ACLs, enabling more flexible access control in IPv6-enabled deployments.

ZooKeeper component upgraded to 3.8.5

ZooKeeper is now upgraded to 3.8.5. This release incorporates stability and correctness fixes from the upstream, which deliver performance and maintenance improvements, and enhanced client connectivity support.

Fixed Issues In Cloudera Runtime 7.3.2

Fixed issues represent issues reported by Cloudera customers that are addressed in this release

Fixed Issues in Atlas

Fixed issues for Atlas are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixed issues in this release.

Cloudera Runtime 7.3.2.10000 SP1

CDPD-99217: Incomplete Metadata Display for Ozone Buckets in Atlas Search Results

7.3.2.10000

While searching for certain Ozone buckets in the Atlas UI, the results page loaded, but key details such as Properties, Classifications, Relationships, and Audits were not displayed. The issue stemmed from the large size of the response payload, specifically the relationships section, which exceeded 500MB in the JSON response. This overwhelmed the UI, preventing it from rendering associated metadata. To optimize the Entity Detail page load time, the GET API now excludes relationship attributes and referred entities from the initial response. Details for the Schema tab, Relationships tab, and Terms are now fetched using a dedicated Relationship API. The Relationships tab now uses a card view with on-scroll pagination, and the Schema tab includes UI caching and pagination support.

CDPD-102546: Rename scenarios are excluded from parallel historical entity resolution

7.3.2.10000

In the distributed notification processing framework, rename scenarios during historical entity resolution are not supported. RENAME events require complex mapping tables and multi-step fallback queries that introduce significant performance overhead. Consequently, lineage adjustments following a rename operation might exhibit discrepancies if the messages are processed out of sequence across parallel topics.

CDPD-98539: Add operation type to the lineage graph

7.3.2.10000

Apache Impala lineage events did not include an explicit operation type, so Apache Atlas inferred the operation from queryText and could miss or misclassify it for some statements when the query spanned multiple lines or started with a SQL comment. Apache Impala lineage events now include an explicit operationType field, and Apache Atlas uses that value when it builds Impala process and column lineage. For queries that insert into a table or overwrite a table, the process segment in lineage qualified names uses INSERT or INSERT_OVERWRITE instead of QUERY or QUER Y_WITH_CLAUSE.

CDPD-103032: Business metadata with an array attribute type is slow to delete

7.3.2.10000

When you deleted a business metadata typedef that included an attribute with the array data type, Apache Atlas could be extremely slow on large catalogs because it scanned Solr for entity references even though array attribute values are not Solr-indexed. Deletion now uses graph-based existence checks instead of Solr scans, which significantly reduces delete time on large catalogs. For business metadata typedef deletes through the REST API, you can use the optional force query parameter on DELETE /api/atlas/v2/types/typedefs and DELETE /api/atlas/ v2/types/typedef/name/{typeName}. When force=true, reference validation is skipped entirely. This parameter applies only to business metadata typedefs, and other typedef categories retain the existing validation behavior.

Apache Jira: ATLAS-4988

CDPD-91502: JanusGraph property keys are created before vertices and edges use them

7.3.2.10000

Atlas could attempt to create graph vertices or edges before the corresponding JanusGraph property keys existed, which caused intermittent graph update failures. The fix ensures property keys are created before they are referenced when vertices and edges are written.

Apache Jira: ATLAS-5073

CDPD-91625: JanusGraph transactions are committed or rolled back reliably

7.3.2.10000

Graph management operations did not always call commit or rollback on AtlasGraphManagement instances, which could leave JanusGraph transactions open and cause instability under load. Each graph management transaction is now closed explicitly.

Apache Jira: ATLAS-5091

CDPD-92009: Hive shell entities can be imported when mandatory attributes are not yet populated

7.3.2.10000

When you import Hive metadata using the import utility in Public Cloud, validation of mandatory attributes for shell entities could fail and block import with an error such as hive_table.name: man datory attribute value missing in type Asset. The fix skips mandatory-attribute validation for shell entities so import can complete and entities can be updated as metadata becomes available.

CDPD-91961: Incremental import handles mixed relationship types under one attribute

7.3.2.10000

When importing entities where a single relationship attribute (such as hive_db.tables) contained a mix of different entity types (for example, both hive_table and iceberg_table references), the import operation could fail or produce incomplete relationship graphs because Atlas incorrectly derived only one relationship type for the entire collection. Atlas now resolves each referenced entity with the correct relationship definition so validation succeeds and all relationships import correctly.

Apache Jira: ATLAS-5156

CDPD-102903: Remote streaming is disabled in Apache Solr configuration for Apache Atlas

7.3.2.10000

Remote streaming in the Apache Solr configuration for Apache Atlas is now disabled by default because it has known security implications. The enableRemoteStreaming="true" attribute is removed from the <requestParsers> element in solrconfig.xml.

CDPD-101880: Indexing is disabled for Hive DDL queryText attributes to reduce write latency

7.3.2.10000

The queryText attribute on the ddl typedef (used by hive_table_ddl and hive_db_ddl) was indexed. Large DDL text values caused high write latency when Hive DDL entities were created or updated. A typedef patch (TYPEDEF_PATCH_0008_001) sets isIndexable to false for ddl.queryText and disables the existing composite index for that attribute during bootstrap. DSL searches that filter on hive_table_ddl.queryText or hive_db_ddl.queryText may take longer than before.

Apache Jira: ATLAS-5229

CDPD-92014: React UI table headers display correctly when no records are present

7.3.2.10000

In the React-based Atlas user interface, table headers in the Search and Business Metadata sections could appear visually constrained or misaligned when a table had zero rows. The table headers now display with proper alignment and layout even when no records are present.

Apache Jira: ATLAS-5151

CDPD-94690: React UI classification and glossary search use the selected item

7.3.2.10000

In the React-based Atlas user interface, choosing Search on a sub-classification incorrectly ran a search for the parent classification instead of the sub-classification you selected; search now runs for the classification you click. Glossary term search from the menu now uses the full hierarchical term name, which returns the correct results. The Create Business Metadata button is aligned in the filter bar, and filter bars are no longer shown on the Classification Detail and Term Detail pages.

Apache Jira: ATLAS-5168

CDPD-95658: React UI displays the isIncomplete property for shell entities

7.3.2.10000

The isIncomplete property was not shown on shell entity detail pages in the React-based Atlas user interface, although it appeared in the Classic UI. Shell entities are placeholder entities created with minimal metadata, and the isIncomplete property indicates that these entities require additional information to be complete. A faulty filter in the property list excluded isIncomplete regardless of value. The property now displays for shell entities (whether true or false), matching Classic UI behavior.

Apache Jira: ATLAS-5170

CDPD-95659: Business Metadata cardinality matches the Enable multiple values setting

7.3.2.10000

When you created business metadata in the Atlas user interface, multi-valued attributes could be saved with SINGLE cardinality even when Enable multiple values was selected. The UI now sets cardinality to SINGLE when Enable multiple values is cleared, to SET by default when it is selected, and to LIST when you toggle the option after creation.

Apache Jira: ATLAS-5169

CDPD-95669: React UI Basic Search entity type links run a search for that type

7.3.2.10000

In the React-based Atlas user interface, clicking an entity type in the type column of basic search results redirected to the home page instead of filtering results to that type. Clicking a type in the column now opens basic search results filtered to that entity type using a DSL query for the type name.

Apache Jira: ATLAS-5172

CDPD-102081: React UI shows Unauthorized when Ranger denies tag or term assignment

7.3.2.10000

When Ranger denied permission to add a classification or glossary term, the React-based Atlas user interface could redirect you to the home page during the action. Atlas now displays an Unauthorized notification from the API response and keeps you on the current page. The import module also shows the full uploaded file name on hover, and the business metadata edit window no longer lets you remove existing attribute types.

Apache Jira: ATLAS-5275

CDPD-99261: searchRelatedEntities supports pagination and returns complete relationship results

7.3.2.10000

When an entity had a very large relationship graph, the searchRelatedEntities API (GET /v2/search/relationship) could return all related entities in one response, which caused slow load times or an unresponsive Atlas user interface. The API now supports pagination, and the Atlas user interface loads the Relationships tab on the entity detail page in pages instead of retrieving the entire graph at once. The API always applied default alphabetical sorting by name even when you did not request sorting, which added overhead for large result sets and slowed deep pagination. The optional disableDefaultSorting query parameter (set to true) skips default name sorting for faster natural graph traversal order; default alphabetical ordering is unchanged when you omit the parameter.

When an attribute can reference multiple entity types (for example, hive_db.tables containing hive_table and hbase_table), the API previously returned entities of only one type and omitted other valid related entities; for example, it could not return complete results when a Hive database contained both Hive and Apache Iceberg tables. searchRelatedEntities now traverses all applicable relationship types and returns every related entity across mixed entity types. Approximate relationship counts could be incorrect when you excluded deleted entities, because the count logic filtered by relationship (edge) status instead of related entity (vertex) status; counting now evaluates entity status and returns accurate totals, including relationships that span multiple entity types.

Apache Jira: ATLAS-5194, ATLAS-5241

CDPD-96774: Entities not created in Atlas when hive table created on top a OFS bucket or volume

7.3.2.10000

When you created a Hive external table with data stored directly on an Ozone File System (OFS) bucket or volume (without a key path under the bucket), Atlas did not create the corresponding entities. Only Hive tables that used OFS key-based paths (for example, ofs://serviceId/volume/bucket/key) were tracked in Atlas; tables that pointed directly to a bucket (for example, ofs://serviceId/volume/bucket) or a volume (for example, ofs://serviceId/volume) did not. Atlas now audits and creates Ozone volume, bucket, and key entities for Hive tables stored on OFS bucket and volume locations.

Apache Jira: ATLAS-5034

CDPD-99096: Perform Degradation: Atlas Replication is 3-4x Slower in 7.3.2 vs. 7.1.9

7.3.2.10000

A performance degradation was observed in the asynchronous table replication process after upgrading from version 7.1.9 to 7.3.2, where the average time to replicate a single table increased significantly. This issue is resolved. The asynchronous import path is optimized by caching intermediate states, consolidating import and bookkeeping operations to reduce graph transactions by more than 50%, and updating the processed entities structure to improve lookup efficiency. These changes significantly improve asynchronous import throughput without altering replication logic.

OPSAPS-75559: Ranger policies secure Kafka topics for Atlas parallel notification processing

7.3.2.10000

When you enable distributed notification processing, Atlas routes metadata and lineage events to Kafka topics with the ATLAS_METADATA_* and ATLAS_LINEAGE_* prefixes. Cloudera Manager adds default Apache Ranger Kafka policies that grant the Atlas service user permission to create, configure, publish, consume, and delete those topics so the atlas-notification-processor service and Atlas server can use them for parallel ingestion.

Cloudera Runtime 7.3.2.100 CHF 1

CDPD-103625: Remote streaming should be disabled as this feature has known security implications.

7.3.2.100

Remote streaming in the Apache Solr configuration for Apache Atlas is now disabled by default because it has known security implications. The enableRemoteStreaming="true" attribute has been removed from the <requestParsers> element in solrconfig.xml.

CDPD-91624: Perform Degradation: Atlas Replication is 3-4x Slower in 7.3.2 vs. 7.1.9

7.3.2.100

A performance degradation was observed in the asynchronous table replication process after upgrading from version 7.1.9 to 7.3.2, where the average time to replicate a single table increased significantly. This issue is resolved. The asynchronous import path is optimized by caching intermediate states, consolidating import and bookkeeping operations to reduce graph transactions by more than 50%, and updating the processed entities structure to improve lookup efficiency. These changes significantly improve asynchronous import throughput without altering replication logic.

CDPD-94458: searchRelatedEntities - Add pagination, optional sorting, fix approximate count while fetching relationships of an entity

7.3.2.100

When an entity had a very large relationship graph, the searchRelatedEntities API (GET /v2/search/relationship) could return all related entities in one response, which caused slow load times or an unresponsive Atlas user interface. The API now supports pagination, and the Atlas user interface loads the Relationships tab on the entity detail page in pages instead of retrieving the entire graph at once. The API always applied default alphabetical sorting by name even when you did not request sorting, which added overhead for large result sets and slowed deep pagination. The optional disableDefaultSorting query parameter (set to true) skips default name sorting for faster natural graph traversal order; default alphabetical ordering is unchanged when you omit the parameter.

When an attribute can reference multiple entity types (for example, hive_db.tables containing hive_table and hbase_table), the API previously returned entities of only one type and omitted other valid related entities; for example, it could not return complete results when a Hive database contained both Hive and Apache Iceberg tables. searchRelatedEntities now traverses all applicable relationship types and returns every related entity across mixed entity types. Approximate relationship counts could be incorrect when you excluded deleted entities, because the count logic filtered by relationship (edge) status instead of related entity (vertex) status; counting now evaluates entity status and returns accurate totals, including relationships that span multiple entity types.

Apache Jira: ATLAS-5194, ATLAS-5241

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Atlas issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

CDPD-82058: UI displays current time instead of "NA" when server returns date fields as '0'

7.3.2

When the Atlas server returns date field values as 0 (zero) in API responses, the user interface incorrectly displays these fields as the current system time instead of showing "NA" (Not Available)

to indicate that no valid date is present. This causes confusion when viewing entity details, audit logs, or other date-related information where timestamps are expected but not available. The fix modifies the UI utility functions in both the classic and React-based user interfaces to properly detect zero date values and display "NA" instead of converting them to the current time.

Apache JIRA: ATLAS-5015

CDPD-84505: Advanced Search not working properly in Atlas UI

7.3.2

When using the Advanced Search feature in the Atlas user interface, selecting a type definition (typeDef) and clicking the search button fails to return any results or fetch entities for the selected type. This issue affects the search result layout view in both the classic and React-based user interfaces. The broken functionality forces users to rely on Basic Search instead of the more powerful Advanced Search capabilities. The fix corrects the search result handling logic to properly process and display search results for type-based queries.

CDPD-87670: Atlas Glossary becomes unresponsive when page size is set to 50

7.3.2

When viewing a glossary item that has more than 25 associated entities and changing the page limit to 50 in the Glossary section, the Atlas user interface becomes unresponsive and displays an error alert with the message "expected type AtlasGlossaryCategory; found AtlasGlossaryTerm". This issue was caused by the UI calling an incorrect API endpoint. The fix corrects the API call to properly handle larger page sizes in the Glossary section.

Apache JIRA: ATLAS-5067

CDPD-82062: UI displays current time instead of "NA" when server returns date fields as '0'

7.3.2

The fix updates date handling utilities in both the classic and React-based Atlas user interfaces so that when API responses contain date field values as 0 (zero), the UI displays "NA" (Not Available) instead of converting them to the current system time.

Apache JIRA: ATLAS-5015

CDPD-89296: Basic search API fails to validate nested and grouped filter criteria

7.3.2

Previously, the /v2/search/basic API endpoint validated only top-level, non-grouped filter criteria. When entity or tag filters used grouped conditions (created with the Add Group Filter option in the Atlas UI), the API did not recursively validate nested FilterCriteria objects. This caused requests with composite or grouped filters to fail with a bad request error such as "Invalid operator specified for attribute: null". The fix adds recursive validation that correctly handles both grouped filter conditions (which require an AND/OR condition and contain nested criteria) and leaf-level filter criteria (which require a valid operator, a non-blank attribute name, and a non-blank attribute value, except when using IS_NULL or NOT_NULL operators). Requests that fail validation now return a descriptive error message.

Apache JIRA: ATLAS-5053

CDPD-81913: Auto-created Kafka topics are not audited to Atlas on ZooKeeper-based clusters

7.3.2

When a Kafka topic was automatically created (for example, following a produce attempt on a non-existent topic), Kafka set the error code to 3 (UNKNOWN_TOPIC_OR_PARTITION) in the MetadataResponse. The Atlas Kafka hook uses an entity selector that only processes audit events with a zero error code, causing topic auto-creation events to be silently dropped and not audited to Atlas. This issue only affected ZooKeeper-based Kafka clusters; KRaft-based clusters handle topic creation differently and were not affected. The fix updates the Atlas Kafka hook to recognize and

correctly audit topic auto-creation events even when the MetadataResponse contains a non-zero error code.

OPSAPS-69156: Use '=' with add-opens/add-modules/add-exports

7.3.2

The JVM option syntax used by Atlas is normalized to use the equals format for module-related flags (for example, --add-opens=..., --add-exports=..., and --add-modules=...). This avoids failures in environments where JAVA_TOOL_OPTIONS does not correctly handle the space-separated variant and improves portability across JDK runtimes.

CDPD-80582: The Atlas client should retry on HTTP 500 errors

7.3.2

When the Atlas plugin attempted to bootstrap the Kafka or Schema Registry model and the Atlas server returned an HTTP 500 error (for example, error code ATLAS-500-00-005 with the message "Failed to get the lock; another type update might be in progress"), the client did not retry the request. This caused the Atlas plugin within Schema Registry to fail initialization permanently, resulting in subsequent errors such as "schema_metadata_info: Unknown/invalid typename". The fix updates the Atlas client retry logic to also retry requests that fail with HTTP 500 (Internal Server Error), in addition to the already-handled HTTP 503 (Service Unavailable) errors.

Apache JIRA: ATLAS-4571

CDPD-69213: Export/Import, Incremental Export: When entity exported has a tag propagated from entity which is deleted, tag is not propagated to it at target

7.3.2

When creating tables with multiple levels of depth, a tag applied to the first table was propagated along the lineage. After dropping the first table and exporting the entire lineage, child tables did not have the propagated tag after import on the target cluster. The root cause was a task ordering issue in the deferred actions flow: when deferred actions were enabled during import, the add-propagation task was created before the parent entity was deleted, but ran after the deletion, finding the edge between the parent entity and the classification in a deleted state and therefore not propagating the tag to child entities. The fix adds a check in the tag propagation logic to skip the edge-status validation when import is in progress and deferred actions are enabled, ensuring that tags are correctly propagated to child entities even when the source entity has been deleted.

Apache JIRA: ATLAS-5055

CDPD-68761: Atlas 'updateTime' parameter is not updated when term is added

7.3.2

When a glossary term was assigned to or removed from an entity, the entity's updateTime metadata field was not updated to reflect the modification. Clients or processes that relied on updateTime to detect entity changes would not recognize updates made through term assignment or removal. The fix calls the entity modification metadata update after each term assignment and dissociation operation in the glossary term processing logic, ensuring that updateTime correctly reflects the latest change.

Apache JIRA: ATLAS-4848

CDPD-67277: Improve efficiency in access request handling by optimizing classification retrieval

7.3.2

When processing bulk classification operations on large numbers of entities, the authorization layer retrieved full classification objects for each entity to perform access checks, even though only the classification names were required. This caused unnecessary graph traversals and redundant data loading, resulting in significantly increased processing time as the number of entities and classifications grew. The fix optimizes the access request handling by using classification names directly instead of iterating over full classification objects, and by replacing the heavier entity header retrieval call that fetched all classification objects with a lighter call during authorization

checks. Performance testing with 260 entities showed processing time reductions of up to 50% depending on the number of classifications involved.

CDPD-79730: Provide liveness and readiness probes in Atlas

7.3.2

Atlas now provides two REST API endpoints to support container orchestration health checks. The liveness probe endpoint (GET /api/atlas/admin/liveness) returns HTTP 200 when the Atlas service state is ACTIVE or MIGRATING, and HTTP 500 when the service requires a restart. The readiness probe endpoint (GET /api/atlas/admin/readiness) returns HTTP 200 only when the Atlas service is ACTIVE and both the index store and the backend store are active, and HTTP 500 otherwise. These endpoints enable container orchestration platforms such as Kubernetes to correctly determine when to restart an Atlas pod and when to route traffic to it.

Apache JIRA: ATLAS-4826

CDPD-70221: Atlas incremental export takes too much time

7.3.2

When performing an incremental export of Hive database entities, Atlas traversed every connected entity in the graph and loaded the full entity object, including all extended information, for each one solely to check whether its modification timestamp fell within the export window. This resulted in data loading overhead equivalent to that of a full bootstrap export, regardless of how few entities had actually changed. The fix introduces an optimized code path for incremental exports of Hive database entities: Atlas now reads the modification timestamp directly from the graph vertex property and loads the full entity only when the timestamp qualifies, avoiding unnecessary data loading for entities that have not been modified since the last export.

Apache JIRA: ATLAS-4979

CDPD-89927: Option to customize name of the configuration filename

7.3.2

Atlas now supports a system property named atlas.properties to specify the configuration filename to load, instead of being limited to atlas-application.properties. This enables hooks and utilities to use custom property filenames (for example, -Datlas.properties=custom-atlas.properties).

Apache JIRA: ATLAS-5098

Fixed Issues in Apache Avro

Review the list of Avro issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixed issues in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new fixed issues in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new fixed issues in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Avro issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

CDPD-88356: Migrate Apache Commons Lang dependency to version 3.18.0 to address CVE-2025-48924

7.3.2

This fix addresses a high-severity uncontrolled recursion vulnerability in Apache Commons Lang's ClassUtils.getClass(…) method by upgrading commons-lang3 dependency to version 3.18.0.

CDPD-85710: Upgrade jQuery for Avro artifacts

7.3.2

Several CVEs (CVE#2020#23064, CVE#2019#5428, NSWG-ECO#328, NSWG-ECO-329) affected the jQuery version used in Avro's web artifacts. jQuery has been upgraded to a secure version address these vulnerabilities.

CDPD-82755: Restrict trusted packages in ReflectData and SpecificData

7.3.2

This fix addresses a deserialization vulnerability (CVE-2024-47561) in ReflectData and Specific Data. The change introduces the org.apache.avro.SERIALIZABLE_PACKAGES system property to restrict class types during serialization, protecting the host system from malicious payloads.

Apache Jira: AVRO-3985

CDPD-77463: Upgrade Avro to version 1.12.0

The Avro component has been upgraded from version 1.11.1 to version 1.12.0.

CDPD-77022: Migrate dependency to commons-lang3

7.3.2

This fix ensures that all the Avro artifacts use the centralized commons-lang3 version.

CDPD-75119: Migrate Avro commons-io dependency to version 2.17.0

7.3.2

This fix addresses tCVE-2024-47554 by upgrading commons-io dependency to version 2.17.0.

Fixed Issues in Cloud Connectors

Review the list of Cloud Connectors issues that are resolved in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Cloud Connectors issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

There are no fixed issues in this release.

Fixed Issues in Cruise Control

Fixed issues and resolved maintenance items for Cruise Control are addressed in Cloudera Runtime 7.3.2 and its associated service packs.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Cruise Control issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

There are no fixed issues in this release.

Fixed issues in Iceberg REST Catalog

Fixed issues for Cloudera Iceberg REST Catalog are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixed issues in this release.

Fixed Issues in HDFS

Fixed issues and resolved maintenance items for HDFS are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.2.10000 SP1

CDPD-108124: HDFS balancer fails to run as a service due to duplicate metrics registration

7.3.2.10000

Previously, when running the HDFS balancer as a long-running service, subsequent balancing iterations would fail with a MetricsException. This occurred because the balancer attempted to register its metrics source multiple times, which is not permitted by the metrics system after initialization. This issue has been resolved by ensuring that balancer metrics are not duplicated across iterations.

Apache JIRA: HDFS-17648

Cloudera Runtime 7.3.2.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves HDFS issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

CDPD-75403: Local information disclosure in RunJar temporary directory

7.3.2

Previously, the temporary directories created by the RunJar.run() command lacked restricted permissions by default. This could allow other local users on the shared system to access sensitive files within those directories. This issue is now resolved. The system now implements stricter access controls on the RunJar working directory, ensuring that only the owner has permission to access the contents.

Apache JIRA: HADOOP-19031

CDPD-69173: FsImage contains inaccessible nodes due to corruption

7.3.2

Previously, corrupted FsImage files could contain inaccessible nodes that the FsImageValidation tool was able to identify but not resolve.. This issue is now fixed. The FsImageValidation tools is enhanced to remove the inaccessible nodes during the validation of the INodeMap if it detects corruption.

Apache JIRA: HDFS-17380

CDPD-88767: OpenSSL 3.x symbols are not loading

7.3.2

Previously, Hadoop native code failed to link with OpenSSL 3.x because specific symbols—EVP_CIPHER_CTX_block_size and EVP_CIPHER_CTX_encrypting—were removed from the library's public API. This issue is now resolved. The native code is updated to use OpenSSL 3.x compatible accessor functions, ensuring proper loading and linking.

Apache JIRA: HADOOP-18583

CDPD-78928: Backport HADOOP-16299 to avoid module access violation in LdapGroupsMapping

7.3.2

Previously, when LdapGroupsMapping was the selected group mapping method, commands relying on this group mapping failed with a Java module access violation on JVM version higher than Java11. This issue is now resolved. The group mapping implementaion avoids illegal access while maintaining the same functionality.

Apache JIRA: HADOOP-16299

CDPD-84571: Backport upstream fixes for missing blocks issue

7.3.2

Backported the following upstream fixes:

Apache JIRA: HDFS-14303, HDFS-15764, HDFS-16985, HDFS-17342

COMPX-20253: Missing Secure Sockets Layer (SSL) cipher inclusion list support

7.3.2

Previously, support for configuring an inclusion list of SSL ciphers for HttpServer2 and SSLFactorywas missing. This issue is now resolved. When an inclusion list is set, only the listed ciphers are allowed, and any cipher present in both the inclusion and exclusion lists are excluded.

Apache Jira: HADOOP-19546

COMPX-23154: Unecessary direct dependency on Jersey 1 and JSR311 APIs fro HTTP headers

7.3.2

Previously, direct dependency on Jersey 1 `jsr311` API for HTTP header handling relied on the javax.ws.rs.core.HttpHeaders class. This issue is now resolved. All usages of javax.ws.rs.core.HttpHeaders are replaces with internal constants or alternative classes, successfully removing the direct dependency.

Apache Jira: HADOOP-19077

Fixed Issues in HBase

Review the list of HBase issues that are resolved in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

CDPD-105335: HBase RPC client connections fail on TLS-enabled clusters

Previously, the HBase RPC client offered only a restricted set of TLS cipher suites when establishing a connection. On clusters with data-in-motion encryption enabled, the client and the HBase RegionServer service could fail to negotiate a common cipher suite, causing TLS handshake errors. As a result, dependent operations failed.

This issue is now fixed. The HBase client now supports a broad range of TLS cipher suites and you can configure and enforce TLS cipher suites on the server side across Master and RegionServer services.

Apache Jira: HBASE-30212

Cloudera Runtime 7.3.2.10000 SP1

CDPD-105033: Missing scanner instance owner check in HBase Thrift and REST services

Previously, HBase Thrift and REST did not verify scanner ownership on fetch and close operations in the scan workflow (open, fetch, close). An authenticated user could read rows from scanners opened by other users or close scanners belonging to other users. For more information, see CVE-2026-49326.

This issue is now fixed. The services verify that the effective user matches the scanner owner before allowing fetch or close operations.

Apache Jira: HBASE-30183

Cloudera Runtime 7.3.2.100 CHF 1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves HBase issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

There are no fixed issues in this release.

Fixed Issues in Hive

Fixed issues for Hive are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

CDPD-121757: Lead and lag window functions returned incorrect results when vectorization was enabled

7.3.2.20000

Previously, the LEAD and LAG window functions returned incorrect results when vectorization was enabled and the queried table contained multiple columns.

This issue is resolved by correcting the vectorized execution logic for the LEAD and LAG window functions so that the results now match non-vectorized execution.

Apache Jira: HIVE-29262

CDPD-121756: Concurrent insert operations on a new dynamic partition caused data loss in non-ACID tables

7.3.2.20000

Previously, when multiple concurrent INSERT INTO operations targeted a newly created dynamic partition in a non-ACID table, some of the inserted data could be lost.

This issue is resolved by preventing concurrent writes to a newly created dynamic partition in non-ACID tables from causing data loss.

Apache Jira: HIVE-29744

Cloudera Runtime 7.3.2.10000 SP1

Hive user's experience

- CDPD-94141: This enhancement displays a clean execution error message, instead of a full stack trace, when you run the `SHOW PARTITIONS` command on a non-partitioned table.

  Apache Jira: HIVE-26926
- CDPD-94142: This enhancement displays a clearer execution error message when you provide an invalid identifier or literal in the STORED BY clause.

  Apache Jira: HIVE-27957
- CDPD-94143: This enhancement labels the status explicitly in the LLAP task scheduler log when a task attempt is registered, which makes the log message easier to interpret.

  Apache Jira: HIVE-28246
- CDPD-94144: This enhancement applies custom user-defined functions (UDFs) registered for vectorized mode during query execution, which expands vectorization coverage across the query plan.

  Apache Jira: HIVE-28830
- CDPD-94145: This fix prevents a NullPointerException in the Tez session pool manager when the active resource plan returns a null plan while logging the applied triggers.

  Apache Jira: HIVE-29007

Hive's performance and optimization

- CDPD-95734: Hive can now reduce the number of rows produced by GROUP BY queries that use grouping sets, which previously could slow down query execution. When the estimated output exceeds a threshold, Hive partitions the input data with an extra shuffle to lower the row count and improve performance. You control this with the new hive.optimize.grouping.set.threshold property, which defaults to 1,000,000,000 rows; set a negative value to turn the optimization off.

  Apache Jira: HIVE-28489
- CDPD-94138: Hive now keeps the original column size when you apply the upper or lower function, so queries run with a more efficient plan.

  Apache Jira: HIVE-28196

- CDPD-94139: Removes excessive logging during hash table flushes and introduces the hive.map.aggr.hash.flush.size.percent property to tune flush sizes and prevent map-side hash aggregation performance degradation

  Apache Jira: HIVE-28428
- CDPD-95735: Hive now estimates a row reduction for IN conditions even when column statistics are unavailable, which produces better query plans.

  Apache Jira: HIVE-28363

Cloudera Runtime 7.3.2.100 CHF 1

CDPD-98415: Hive ACID incremental replication no longer fails with "MetaException: invalid write id"

7.3.2.100

Previously, during Hive ACID incremental replication involving large volumes of events (such as bulk DDLs, partition updates, or TRUNCATE TABLE operations), the replication load task failed with the following error:

```
Caused by: org.apache.hadoop.hive.metastore.api.MetaException: i
nvalid write id <ID> for table <db_name>.<table_name>
```

This occurred because the background task responsible for cleaning up dangling transactions (CLEAR_DANGLING_TXNS) executed prematurely at the end of individual task batch iterations instead of waiting for the full incremental load operation to complete. As a result, the cleanup process aborted in-flight target transactions before replaying their corresponding commit events in subsequent batches, which invalidated the target write ID.

Cloudera Manager and Cloudera Runtime now defer the CLEAR_DANGLING_TXNS task to execute only after all task iterations in the incremental load operation complete successfully. This fix prevents premature transaction aborts and ensures seamless incremental replication.

Apache Jira: HIVE-29459

CDPD-103320: Cleaner deletes base directories incorrectly during ACID table cleanup

7.3.2.100

Previously, when the minimum open write ID was higher than the cleaner high-water mark, the Hive ACID cleaner utilized an incorrect shortcut that skipped directory validation.

This issue is resolved by removing the incorrect directory validation shortcut. The cleaner now checks all base directories against the high-water mark to prevent premature deletions and protect data integrity.

Apache Jira: HIVE-29601

CDPD-104581: Queries return incorrect results when combining full outer join, map join, and filter expressions

7.3.2.100

Previously, Hive produced incorrect query results when processing a full outer join combined with a map join and filter expressions.

This issue is resolved by updating the query optimizer logic to prevent the conversion of a full outer join to a map join when filter expressions are present.

Apache Jira: HIVE-27988

CDPD-104580: Database inconsistency during partition deletion failures

7.3.2.100

Previously, when a direct SQL processing failure occurred during operations such as dropping partitions, the partial changes made to the database were not reverted before the system attempted an alternative processing method.

This issue is now resolved. The system now sets a savepoint before starting direct SQL processing, ensuring that if a failure occurs, the transaction rolls back to the savepoint before retrying the operation.

Apache Jira: HIVE-26976

CDPD-101750: Compaction retries causing potential data loss

7.3.2.100

Previously, compaction retries triggered by worker timeouts under an incorrect configuration caused a risk of data loss. If a compaction was re-attempted before the prior attempt completed or was properly aborted, the Cleaner observed multiple base directories with the same write ID and mistakenly deleted all of them.

This issue is resolved by validating open transactions during worker timeouts, correcting the eligibility conditions for cleaning, and raising exceptions rather than assertions in the transaction utilities to safeguard valid base directories.

Apache Jira: HIVE-29420

CDPD-101699: Deserialization failure when cost-based optimization is disabled

7.3.2.100

Previously, when Cost-Based Optimization (CBO) was disabled, some queries might lead to a plan containing a generic function expression referring to a Guava collection.

This issue is now resolved by replacing the Guava collection with a standard Java list.

Apache Jira: HIVE-29488

CDPD-95663: Incorrect query results during full outer join optimization with filters

7.3.2.100

Previously, when a query contained a full outer join along with filter expressions, optimizing the operation into a map join caused the execution engine to ignore the filters or return duplicated rows.

This issue is now resolved by preventing the conversion of a full outer join to a map join when filter expressions exist.

Apache Jira: HIVE-27988

CDPD-104583: Data race conditions during concurrent partition column statistics updates

7.3.2.100

Previously, updating partition column statistics followed a sequential process of fetching, modifying, and then updating the statistics.

This issue is now resolved by implementing a transactional mechanism that synchronizes concurrent updates securely.

Apache Jira: HIVE-29316

CDPD-104582: Constraint violation errors during concurrent table column statistics replication

7.3.2.100

Previously, when High Availability (HA) was enabled and multiple Hive Metastore clients processed statistics replication across different engines simultaneously, a data race condition caused duplicate key errors and constraint violations.

This issue is now resolved by automatically coordinating the clients so that an update is performed instead of an insert if the data was already saved by a parallel process.

Apache Jira: HIVE-28578

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Hive issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

CDPD-97246: Avro schema literal logging at INFO level

7.3.2

Previously, the Avro deserializer logged the schema literal at the INFO level.

This issue is now resolved by changing the log level to DEBUG.

Apache Jira: HIVE-22606

CDPD-96649: Incorrect aggregate statistics when direct SQL batch retrieval is enabled

7.3.2

Previously, when hive.metastore.direct.sql.batch.size config was greater than 0, the system failed to merge column statistics correctly if the number of partitions or columns exceeded that batch size.

This issue is now resolved by ensuring the statistics are properly merged during the retrieval process, preventing redundant entries.

Apache Jira: HIVE-29203

CDPD-95834: Hive Metastore backend database schema out of date

7.3.2

Previously, the Hive Metastore (HMS) backend database schema files required updates for compatibility with the latest version.

This issue is now resolved by updating the schema files to version 7.3.2.0 to ensure a successful upgrade.

CDPD-95681: Hive Beeline connection failure due to SSL certificate hostname mismatch

7.3.2

Previously, the Beeline client failed to establish a connection to HiveServer2 in ZooKeeper High Availability (HA) setups.

This issue is now resolved by ensuring that the SSL certificates correctly match the DNS hostnames, allowing the secure connection to be verified and established.

CDPD-93766: Missing results during anti-join conversion

7.3.2

Previously, queries involving specific join conditions returned empty results instead of the expected data when anti-join conversion was enabled.

This issue is now resolved by preventing anti-join conversion in these specific scenarios to ensure query result accuracy.

Apache Jira: HIVE-29175

CDPD-93756: SHOW COMPACTIONS output filtering

7.3.2

Previously, the `SHOW COMPACTIONS` output included all historical information, which caused the display to become unwieldy when many partitions or history lines existed.

This issue is resolved by adding the ability to filter the command output by database, table, partition, and compaction type or state.

Apache Jira: HIVE-13353

CDPD-93617: Duplicate records during minor compaction

7.3.2

Previously, when the Hive Metastore (HMS) crashed, active compaction jobs were incorrectly reset.

This issue is now fixed by updating the compactor cleaner to address duplicate directories.

Apache Jira: HIVE-29210

CDPD-93432: Incorrect results for anti-join queries

7.3.2

Previously, the HiveAntiJoin rule incorrectly replaced IS NULL filters on nullable columns, which resulted in missing records in the query output.

This issue is now fixed by improving the logic within the HiveAntiSemiJoinRule to ensure accurate query plans.

Apache Jira: HIVE-29176

CDPD-93431: Incorrect results for n-way joins

7.3.2

Previously, queries produced incorrect results when an n-way join contained a combination of both anti and outer joins.

This issue is now resolved by extending the CommonJoinOperator to properly support anti joins when they are used alongside outer joins in n-way join operations.

Apache Jira: HIVE-29290

CDPD-93428: Manifest files appearing in read queries

7.3.2

Previously, temporary directories containing direct insert manifest files used a prefix that allowed them to be included in concurrent read queries.

This issue is now resolved by hiding the direct insert manifest directory from read queries, ensuring only valid data files are processed.

Apache Jira: HIVE-29297

CDPD-93425: Data loss during minor compaction

7.3.2

Previously, query-based minor compaction incorrectly used the minimum open write ID as a lower bound for selecting data files.

This issue is now fixed by removing the incorrect check against the minimum open write ID, allowing the high-watermark to correctly define the compaction range.

Apache Jira: HIVE-29272

CDPD-92923 / CDPD-92586: Memory leak in Hive Metastore REST Catalog

7.3.2

Previously, the Hive Metastore (HMS) REST Catalog leaked memory during `ALTER TABLE` authorization checks. Each check created new catalog instances and handler pools with JMX enabled, which prevented the system from reclaiming memory.

This issue is now fixed by reusing existing catalog instances and disabling JMX for the handler pool to allow the system to recover resources.

CDPD-92499: Hive LDAP authentication and ZooKeeper connections

7.3.2

Previously, Hive Lightweight Directory Access Protocol (LDAP) authentication failed when connecting to a SASL-enforced ZooKeeper instance.

This issue is now resolved by updating the service components to support simultaneous authentication configurations and ensuring correct credential handling during service discovery.

Apache Jira: HIVE-29138

CDPD-92478: Timestamp processing in MetaStoreUtils

7.3.2

Previously, Hive Metastore utilities used local time zone settings to convert between timestamps and strings.

This issue is now fixed by using UTC time zone and the java.time.Instant class to process timestamps, which ensures that time points are represented accurately regardless of local time zone rules.

Apache Jira: HIVE-28337

CDPD-91415: WebHCat and Python script compatibility with Python 3

7.3.2

Previously, hcat.py and various Python scripts used in q files contained syntax that was incompatible with Python 3.

This issue is now resolved by updating the Python scripts to use Python 3-compatible syntax.

Apache Jira: HIVE-25817

CDPD-90670: Incorrect results for queries with multiple lateral view operations

7.3.2

Previously, queries that used two or more `LATERAL VIEW` explode operations along with a `WHERE` clause returned incorrect results when Cost-Based Optimization (CBO) was enabled.

This issue is now resolved by updating the logic to correctly identify separate table aliases for lateral view columns, ensuring that filters are applied accurately.

Apache Jira: HIVE-29084

CDPD-90303: Incorrect results from a CASE expression

7.3.2

A query that used a CASE expression to conditionally return values produced an incorrect result. The query plan incorrectly folded the *CASE statement* into a COALESCE function, which led to a logic error that filtered out some of the expected results.

This issue is addressed by adding a more strict check when converting CASE expressions into COALESCE during query optimization.

Apache Jira: HIVE-24902

CDPD-89462: Performance degradation for wide tables in DirectSqlUpdatePart

7.3.2

Previously, updating or inserting partition statistics for tables with a high number of columns and partitions was slow.

This issue is now resolved by improving the hashing logic to ensure faster data retrieval and insertion, even for tables with thousands of columns and partitions.

Apache Jira: HIVE-29165

CDPD-88987: Improved performance for adding columns to partitioned tables

7.3.2

Previously, metadata operations for adding columns to tables with a high number of partitions and columns were slow because the system utilized a less optimized implementation for partition updates.

This issue is addressed by implementing a more efficient batch processing method for partition updates, which utilizes optimized metadata queries to improve performance for tables with many partitions.

Apache Jira: HIVE-28956

CDPD-88981: Performance degradation during column addition with cascade

7.3.2

Previously, adding columns to a table using the `CASCADE` command resulted in slower performance after optimizations for metadata storage were enabled.

This issue is now resolved. The fix includes an optimized method to reuse column descriptors across partitions, which restores performance levels during the column addition process.

Apache Jira: HIVE-29042

CDPD-88166: Query failure during JDBC filter optimization

7.3.2

Previously, certain queries failed with a class-cast error during the optimization phase.

This issue is now resolved by updating the query optimization rules to ensure that relational operators are correctly identified and processed.

Apache Jira: HIVE-25356

CDPD-87266: Query failure during Tez execution

7.3.2

Previously, Hive queries failed during execution after an upgrade. This resulted in a vertex failure and prevented queries from completing successfully.

This issue is now resolved by updating the execution engine to correctly instantiate internal query split generators during the initialization process.

CDPD-84149: MariaDB connector recognition failure

7.3.2

Previously, Hive failed to recognize the MariaDB connector even when the driver was present.

This issue is now fixed.

CDPD-83461: Query failure when using stack function with union operations

7.3.2

Previously, queries utilizing the `STACK` function in combination with `UNION` operations failed with an internal error during the compilation phase.

This issue is now resolved by updating the query optimization logic to correctly handle the stack function during union operations, preventing the internal processing error.

Apache Jira: HIVE-29029

CDPD-83334: Improving performance for alter partition operations

7.3.2

When altering partitions, the system used Java Data Objects (JDO) updates, which required fetching all fields of old partitions and produced redundant queries, leading to slower performance.

This issue is resolved by implementing direct SQL for altering partitions.

Apache Jira: HIVE-27530

Direct SQL failure during partition alterations

7.3.2

Previously, direct SQL for partition alterations failed in certain database environments due to Character Large Object (CLOB) casting errors and missing boolean type conversion checks.

This issue is resolved by updating the direct SQL logic to handle CLOB types correctly and ensuring proper boolean type conversions during batch updates.

Apache Jira: HIVE-28271

CDPD-80146: Group by alias query failures

7.3.2

Previously, the hive.runtime.dialect.enable property was enabled by default, which caused the hive.groupby.position.alias property to be ignored.

This issue is resolved by setting the hive.runtime.dialect.enable property to false by default. Hive now correctly respects the hive.groupby.position.alias configuration.

CDPD-79144: Incorrect schema version in Hive schema initialization script for MySQL

7.3.2

Cluster creation with 7.3.1 fails due to an incorrect database schema in the CDH_VERSION table.

The issue was addressed by correcting the schema version in the Hive schema initialization script, ensuring successful cluster creation.

CDPD-78337: Merge task not invoked for external CTAS queries on object stores

7.3.2

Previously, the merge task was not invoked for external `Create Table As Select` (CTAS) queries when using S3 or other object stores.

This issue is now resolved by ensuring the merge task is correctly invoked after optimization for external CTAS queries.

Apache Jira: HIVE-27536

CDPD-78329: HiveServer2 runs out of memory with multiple parallel queries with fetch task

7.3.2

Previously, HiveServer2 may run out of memory when multiple parallel queries use fetch task caching (hive.fetch.task.caching=true). This causes queries to fail and HiveServer2 to crash.

The issue was addressed by reducing the default value of hive.fetch.task.conversion.threshold from one GB to 200MB, preventing excessive memory usage and improving stability.

CDPD-77869: Iceberg table data written to HDFS instead of S3 in RAZ-enabled clusters

7.3.2

Previously, when you configured a cluster with Ranger Remote Authorization Board (RAZ) and updated configurations to use S3, data for Iceberg tables was unexpectedly written to the HDFS external table location. This occurred because Iceberg tables, which are treated as external tables, defaulted to the database LOCATION property that still pointed to HDFS if the database was created prior to the S3 switch.

This issue is now fixed by ensuring that external tables correctly align with the intended S3 storage paths. You can now update the database location to point to the S3 bucket to ensure all future external tables default to the cloud storage.

CDPD-75665: Importing a table generates a DDL with an incorrect location

7.3.2

When creating a table using the IMPORT command, this table's partitions could point to an incorrect location, that is, an external imported table can have his partitions located under the managed warehouse directory, which violates the metastore.warehouse.external.dir and metastore.warehouse.dir that intend to host different type of tables.

Apache Jira: HIVE-28580

Hive query execution failure due to AM container exit on lost node with Exit code -100

7.3.2

Hive query failed when ApplicationMaster container was lost

Previously, when running a Hive query, a failed ApplicationMaster (AM) container did not trigger a DAG retry and caused the query execution to fail if the failure message included diagnostic information with a line break.

This issue is now resolved by automatically re-executing the DAG if the AM fails.

Apache Jira: HIVE-28093

CDPD-74539: MariaDB falls back to MySQL in Hive

7.3.2

Hive downstream had errors in supporting MariaDB.

The issue was addressed by making MariaDB automatically fall back to MySQL.

CDPD-66731: Hive Metastore query failure during Zero Downtime Upgrade

7.3.2

Previously, during a Zero Downtime Upgrade (ZDU) from version 7.2.17 to 7.2.18, long-running queries such as `INSERT INTO` statements failed with a MetaException.

This issue is now fixed by ensuring that transaction blocks are correctly managed during the statistics update task, preventing the "current transaction is aborted" error.

CDPD-60770: Passwords with special characters fail to connect with Beeline

7.3.2

When you used a password containing special characters like #, ^, or ; in a JDBC URL for a Beeline connection, the connection failed with a 401 error. This happened because Beeline did not correctly interpret these special characters in the password.

This issue is resolved by introducing a new method to reparse the password from the original JDBC URL, allowing Beeline to correctly handle and authenticate passwords containing special characters.

Apache Jira: HIVE-28805

CDPD-58428: Bucket Map Join hangs when source vertex parallelism changes

7.3.2

Previously, a Bucket Map Join could hang if the parallelism of a source vertex was modified by automated reducer parallelism.

This issue is now fixed by disabling automated reducer parallelism for vertices that serve as a source for a Bucket Map Join.

Apache Jira: HIVE-27078

CDPD-58428: Incorrect results in map-side Sort-Merge Bucket Join with different bucket sizes

7.3.2

Previously, map-side Sort-Merge Bucket (SMB) Joins returned incorrect results when joining two tables with different bucket counts (for example, joining a table with two buckets to a table with three buckets).

This issue is now fixed by implementing a new routing algorithm that ensures bucket N of a small table is correctly mapped to bucket M of a large table based on the greatest common divisor of their bucket sizes.

Apache Jira: HIVE-27357

CDPD-50060: Configurable filter for partition metadata properties in Hive Metastore

7.3.2

Previously, Hive Metastore (HMS) API calls failed with a TTransportException (MaxMessageSize reached) when processing tables with large partition metadata.

This issue is now fixed by providing a configurable filter that excludes unnecessary properties from listPartitions API responses. This change reduces the metadata payload size, prevents connection timeouts, and improves the performance of metadata operations.

Apache Jira: HIVE-27114

CDPD-44551: Avro table import or download fails with ODBC driver due to missing property

7.3.2

The absence of metastore.storage.schema.reader.impl caused Avro table import or download failures in Cloudera runtime 7.1.7 when using the ODBC driver.

The issue was addressed by setting metastore.storage.schema.reader.impl to org.apache.hadoop.hive.metastore.SerDeStorageSchemaReader by default.

Apache Jira: HIVE-26952

CDPD-92208: Query failure when selecting data from views with bracketed definitions

7.3.2

Previously, a `SELECT` query against a view failed with a SemanticException if the view was created with specific column names and a definition enclosed in brackets.

This issue is now fixed by ensuring that the compiler does not add extra brackets if the view definition is already enclosed.

Apache Jira: HIVE-26493

CDPD-83530: Task commits were allowed despite an exception being thrown in the Tez processor

7.3.2

A communication failure between the coordinator and executor caused a running task to terminate, resulting in a java.lang.InterruptedException being thrown by the ReduceRecordProcessor.init(). Despite this exception, the process still allowed the task to be committed and generated a commit manifest.

This issue has now been resolved. The fix ensures that outputs are not committed if an exception is thrown in the Tez processor.

Apache Jira: HIVE-28962

CDPD-89414: Incorrect results for window functions with IGNORE NULLS

7.3.2

When you used the FIRST_VALUE and LAST_VALUE window functions with the IGNORE NULLS clause while vectorization was enabled, the results were incorrect. This occurred because the vectorized execution engine did not properly handle the IGNORE NULLS setting for these functions.

This issue is addressed by modifying the vectorized processing for FIRST_VALUE and LAST_VALUE to correctly respect the IGNORE NULLS clause, ensuring the same results are produced whether vectorization is enabled or disabled.

Apache Jira: HIVE-29122

CDPD-85600: Select queries with ORDER BY fail due to compression error

7.3.2

When you ran a Hive `SELECT` query with an `ORDER BY` clause, it failed with a java.io.IOException and java.lang.UnsatisfiedLinkError related to the zlib decompressor.

The issue was addressed by ensuring the zlib native library is correctly loaded.

Apache Jira: HIVE-28805

CDPD-90301: Stack overflow error from queries with OR and MIN filters

7.3.2

Queries, cause a stack overflow error when they contained multiple OR conditions on the same expression, such as MINUTE(date_) = 2 OR MINUTE(date_) = 10.

This issue is addressed by modifying the HivePointLookupOptimizerRule to keep the original order of expressions and to check if a merge can be performed before creating a new expression.

Apache Jira: HIVE-29208

DWX-20754: Invalid column reference in lateral view queries

7.3.2

The virtual column BLOCK__OFFSET__INSIDE__FILE fails to be correctly referenced in queries using lateral views, resulting in the error:

```
FAILED: SemanticException Line 0:-1 Invalid column reference 'BL
OCK_OFFSET_INSIDE_FILE.
```

This issue is now resolved.

Apache Jira:HIVE-28938

Fixed Issues in Cloudera Data Explorer (Hue)

Fixed issues for Cloudera Data Explorer (Hue) are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.10000 SP1

CDPD-97016: Incorrect table values during CSV import in Data Explorer

7.3.2.10000

Previously, importing CSV files to create Hive tables contained quoted fields with commas through the Data Explorer Importer caused field parsing errors. This issue resulted in malformed table values and corrupted data ingestion. This issue is now resolved, by updating the SQL table creation logic to handle quoted CSV fields correctly and applying the appropriate storage format during import.

CDPD-95672: Directory rename failure in Data Explorer

7.3.2.10000

Previously, attempting to rename a directory in Data Explorer by using the Data Explorer user interface failed, which caused the directory to disappear from the user interface. This issue is now fixed by updating the path handling and user interface feedback during directory renames, ensuring that directories are renamed correctly and remain visible.

CDPD-98406: Excessive Gunicorn debug log generation

7.3.2.10000

Previously, Gunicorn log listeners generated debug-level records for access and metrics even when Data Explorer debug mode was disabled. This behavior caused unnecessary log clutter. This issue is now fixed by ensuring Gunicorn log output honors the Data Explorer debug mode setting.

CDPD-90751: Reverse tabnabbing security vulnerability in external links

7.3.2.10000

Previously, external links using the target="_blank" link attribute without the rel="noopener" or rel="noreferrer" property values created a reverse tabnabbing security vulnerability. This issue is now fixed.

CDPD-101177: Data Explorer File upload failure to S3 in Data Explorer

7.3.2.10000

Previously, file uploads to AWS S3 through Data Explorer might fail in both RAZ and non-RAZ environments because the storage connector incorrectly attempted to use a legacy configuration path. This issue is resolved by ensuring the storage connector uses the correct configuration path for the boto3 SDK.

CDPD-101696: Query Processor failure with Oracle database

7.3.2.10000

Previously, retrieving queries when using an Oracle database failed due to a parameter binding issue in the Query Processor. This issue is now fixed by correcting the parameter binding logic during query retrieval.

CDPD-45130: Truncating excessive length queries to prevent database indexing errors

7.3.2.10000

Previously, large or complex SQL statements, such as lengthy INSERT queries, were indexed by the Query Processor. This resulted in increased load times for the Job Browser. You can now configure query truncation by using the hue.query-processor.query.max-length property in the Query Processor configuration under dasConf section. By default, no truncation is performed to ensure backward compatibility.

CDPD-95666: Persistent loading indicator for unpermitted directories

7.3.2.10000

Previously, selecting a folder or file without read permissions in Microsoft Azure environments caused a persistent loading bar to appear across file systems, including Amazon S3 and HDFS. The user interface remained in a loading state and failed to display an alert message indicating missing permissions. This issue is now fixed.

Cloudera Runtime 7.3.2.100 CHF 1

CDPD-104785: Arbitrary file read on Data Explorer hosts through the localfile importer feature

7.3.2.100

Previously, a local file inclusion vulnerability existed on the Data Explorer host type through the localfile importer feature. This occurred because two indexer API endpoints accepted unvalidated paths, allowing unauthorized file reads through the path parameter. This issue is now fixed.

CDPD-98620: Azure Blob File System upload authentication failure

7.3.2.100

Previously, you experienced intermittent file upload failures by using the Data Explorer user interface. This issue occurred because the token signature parameter value corrupted by collapsing consecutive slashes. As a result, Azure Blob File System storage rejected the authentication request with a 403 Client Error message. This issue is now fixed.

CDPD-91439: Support for Python-OracleDB connectivity in Data Explorer

7.3.2.100

Previously, Data Explorer experienced database connectivity incompatibilities due to limitations with the deprecated cx_Oracle driver. This issue is now fixed by transitioning Data Explorer to use the new python-oracledb library. This change enables official support for modern Python-OracleDB

connectivity, ensures compatibility with newer versions like Python 3.11, and secures the Oracle database integration within Data Explorer.

CDPD-92510: Inaccurate table update timestamps in Data Explorer

7.3.2.100

Previously, the "Data last updated on" timestamp for Iceberg tables in Data Explorer displayed incorrect information when tables were altered by using Spark. This occurred because the timestamp was sourced from the transient_lastDdlTime table property, which does not update correctly during Spark operations. This issue is now fixed by updating the metadata source to use Iceberg snapshot timestamps.

CDPD-92857: Export all failures in secured multi-tenant clusters

7.3.2.100

Previously, the export all feature in Data Explorer failed in secured multi-tenant clusters that used tenant colocation policies. This resulted in access errors during the export process. This issue is now resolved. The export all functionality now follows tenant boundaries and cluster security policies.

CDPD-68779: Error while browsing S3 buckets or ADLS containers from the left-assist panel

7.3.2.100

Previously, when attempting to browse S3 buckets or ADLS containers from the left-assist panel in Data Explorer without the required permissions, the system displayed a generic error message: Failed to retrieve buckets:1:0: syntax error. This is now fixed. When RAZ is enabled, the left-assist panel now opens the home directory if the required permissions are available. In non-RAZ environments, the left-assist panel opens the path configured in the Data Explorer configuration file.

CDPD-100816: Data Explorer File Browser workflow submission button failure

7.3.2.100

Previously, the Submit button component in the Data Explorer File Browser application failed to function, which prevented submitting workflows. This issue is now fixed by restoring the required path reference dependency.

CDPD-96460: Impala query profile row count display error

7.3.2.100

Previously, the Data Explorer Impala editor component displayed incorrect row count values for large tables because it incorrectly parsed the RowsReturned metric property from the EXCHANGE node component inside the Impala query profile log. This issue is fixed by implementing a threshold-based bucketing approach that displays approximate row count indicators, such as 1000 rows or 10000+ rows, to prevent inaccurate query metadata displays.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Impala issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

CDPD-73189: Not all tables are loading in the left assist panel

7.3.2

Previously, in the Data Explorer left assist panel, only the initial 5,000 tables were loaded because the default value of max_catalog_sql_entries was set to 5000. Increasing this configuration in the Cloudera Data Explorer (Hue) Safety Valve only allowed loading up to 10,000 tables due to the max_rows parameter being hardcoded to 10000. This issue is now resolved by fixing the hardcoded limitation. You can now load more than 10,000 tables in the left assist panel.

CDPD-71098: Complex Hive or Impala data type sampling fails in Data Explorer

7.3.2

Previously, sampling complex Hive or Impala data types in Data Explorer did not work correctly. This issue is now fixed. The resolution enables accurate sampling and autocomplete functionality for nested complex types, which improves the query experience.

CDPD-48151: Manual table creation using the importer

7.3.2

Previously, a JavaScript error occurred when you used the Manually option in the importer to create a table. This issue prevented fields from being displayed. This issue is now fixed. To create a table without importing data, you can also use the SQL editor.

CDPD-47230: Missing SQL information in Table Browser for Hive views

7.3.2

Previously, the viewSQL feature in the Table Browser displayed an empty window for views created by using Hive. This issue occurred when the source type was set to Hive, while the feature worked correctly for Impala. This issue is now fixed. You can now view the SQL information for Hive views in the Table Browser as expected.

CDPD-92829: Job Browser error and traceback exposure

7.3.2

Previously, Job Browser displayed an error if the POST format parameter contained empty or invalid values. This issue caused a JSON serialization failure and exposed sensitive system information. This issue is now fixed. The resolution includes validated parameter inputs and improved error handling to prevent the exposure of sensitive system information.

CDPD-77831: Workflow log icon fails to redirect in Data Explorer

7.3.2

Previously, the log icon in the Task tab for workflows did not redirect to the logs. When you clicked the icon during a workflow run, the UI remained on the same page without displaying an error, although a TypeError related to lodash.trimRight occurred in the background. This issue is now fixed. You can now access workflow logs through the UI as expected.

CDPD-74389: File and folder creation fails in Data Explorer S3 file browser

7.3.2

Previously, you could not create files or folders in the Data Explorer S3 File Browser on the AWS-Fedramp-RAZ environment. These actions failed to complete without displaying an error message in the UI. This issue is now fixed.

CDPD-90906: Database connection errors in metrics module

7.3.2

Previously, the metrics module encountered the cx_Oracle.DatabaseError: DPI-1010: disconnect error when the Cloudera Manager metrics API queried active users. This issue caused database connections to be lost. This issue is now fixed. The resolution improves connection stability and robustness against database disconnect errors.

CDPD-89478: Oozie UI displays start and end times incorrectly after upgrade

7.3.2

Previously, the Oozie UI intermittently displayed the start time and end time fields in reverse order. Additionally, the Run button remained inactive until you clicked Save, even if you made no changes. This issue is resolved by correcting the field display order and ensuring that the Run button is active for schedules without requiring an unnecessary save action.

CDPD-58142: Query pre-population fail in the SQL editor

7.3.2

Previously, when you clicked Re-Execute to rerun a query from the Job Browser Queries Query Details page, the query did not appear in the SQL editor as expected. This issue occurred because

the editor loaded multiple times, overwriting the query text. Additionally, the page change was not reflected in the URL. This issue is now fixed. The SQL editor now correctly populates and retains the query text when you navigate from the Job Browser.

CDPD-65749: Table Browser navigation error with Knox

7.3.2

Previously, clicking the back button when navigating the Table Browser directory structure through Knox resulted in a 404 Error. This issue is now fixed. The resolution improves navigation stability and ensures the UI correctly displays the previous page.

CDPD-75031: Impala query status display failure in Data Explorer

7.3.2

Previously, successful Impala queries were displayed as failed in the Data Explorer query history. This issue occurred because Data Explorer did not recognize updated exception message text returned by Impala for expired query handles in upgraded cluster environments. This issue is resolved by updating exception message handling in Data Explorer to recognize revised Impala error strings across cluster versions.

OPSAPS-73942: Query processor service configuration failure during upgrade

7.3.2

Previously, the upgrade wizard failed during the configuration validation phase if the Query processor service was installed. This issue occurred because the service version could not be upgraded directly, resulting in a warning that the service must be removed. This issue is now resolved.

CDPD-81025: Performance issues when querying large datasets with Oracle or MySQL queries

7.3.2

Previously, when querying large datasets using Oracle or MySQL backends, the QueryProcessor failed to correctly apply the SQL limit clause. This caused the system to fetch all data records instead of the restricted set required for pagination. As a result, the user interface experienced significant loading delays and performance issues when displaying query results. This issue is now resolved.

Fixed Issues in Impala

Fixed issues for Impala are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

Impala Kubernetes operator required overly broad RBAC permissions

7.3.2.20000

Previously, the Impala Kubernetes operator was deployed with broad, cluster-admin style role-based access control (RBAC) permissions to simplify initial bootstrap. This level of access was more permissive than production environments required.

This issue is resolved by scoping the operator RBAC permissions to only the API groups and resources required for Helm chart and operator reconciliation, following the principle of least privilege. Core, LDAP, Kudu, and Ranger-enabled deployments continue to reconcile correctly with the reduced permission set.

Apache Jira: IMPALA-15049

Add secure-cluster configuration support to Impala Helm chart

7.3.2.20000

Previously, the Impala Helm chart provided only non-secure default configuration and did not offer a documented way to configure secure-cluster settings, such as TLS and Kerberos, for Impala components deployed on Kubernetes.

This improvement adds structured security-related configuration values and template wiring to the Impala Helm chart, so that TLS and Kerberos-related daemon flags, environment variables, volume mounts, and secrets references can be configured for secure deployments. These security settings remain disabled by default, so existing non-secure quickstart deployments are unaffected.

Apache Jira: IMPALA-15017

Add OAuth authentication options to Impala Helm chart

7.3.2.20000

Previously, the Impala Helm chart supported LDAP and Ranger authentication options but did not expose OAuth authentication configuration for impalad, even though OAuth is commonly used in Kubernetes environments.

This improvement adds OAuth configuration keys to the Impala Helm chart values, so that OAuth authentication can be enabled and configured for impalad in a typed and discoverable way. OAuth support is disabled by default, and the corresponding impalad startup arguments are rendered only when it is explicitly enabled, so existing LDAP-based deployments are unaffected.

Apache Jira: IMPALA-15016

Add typed operator config keys for daemon flags and query defaults

7.3.2.20000

Previously, updating Impala daemon flags and query defaults through the Kubernetes operator required raw, indexed extraArgs values or the general-purpose spec.set field, which was error-prone and difficult to operate.

This improvement adds typed spec.config fields for daemon flags and impalad query defaults, which the operator translates into the corresponding Helm values during reconciliation. The existing spec.set field remains supported as a fallback, so existing custom resources continue to work unchanged.

Apache Jira: IMPALA-14983

Add operation type to the lineage graph

7.3.2.20000

Previously, the lineage event log produced by Impala did not include the operation type of the logged statement. Tools such as Apache Atlas had to infer the operation type by applying regular expressions to the query text, an approach that could not reliably determine the operation type for all statements, for example when the SQL text contained a single-line comment.

This improvement adds an operation type field to the lineage graph that Impala produces, so that lineage consumers no longer need to infer the operation type from the query text.

Apache Jira: IMPALA-14768

Impala returned incorrect results for queries with a predicate on a window function result in an inline view

7.3.2.20000

Previously, Impala could incorrectly migrate an inferred predicate into an inline view when the inline view contained a window function. As a result, the predicate was evaluated against the underlying scan before the window function was applied instead of after it, and queries that combined a window function filter with an additional predicate on the windowed result set could return incorrect results.

This issue is resolved by preventing predicates that depend on window function results from being incorrectly pushed down to the underlying scan, preserving correct results for join conditions and filtered window function queries.

Apache Jira: IMPALA-13262

Cloudera Runtime 7.3.2.10000 SP1

CDPD-105759: Incorrect results from a LEFT JOIN converted to an INNER JOIN

7.3.2.10000

Previously, when the enable_outer_join_to_inner_transformation option was enabled, a `LEFT JOIN` query whose `WHERE` clause used a NOT(... AND ...) predicate was incorrectly converted to an `INNER JOIN`, which dropped rows that contained NULL values and returned incorrect results.

This issue is resolved by correctly identifying nullable predicates so that the join is not converted when the conversion would change the result.

Apache Jira: IMPALA-14685

CDPD-97102: Incorrect matching for a LIKE pattern with a leading wildcard and an escaped trailing percent sign

7.3.2.10000

Previously, a `LIKE` pattern that combined a leading wildcard with an escaped trailing percent sign, such as LIKE '%abc\%', matched any string containing abc\ instead of strings ending in a literal abc%, returning incorrect results.

This issue is resolved by correctly handling the escaped percent sign in the pattern.

Apache Jira: IMPALA-14731

CDPD-98998: Query hangs when the select list contains an unaliased expression with Unicode letters

7.3.2.10000

Previously, a query hung indefinitely, and could not be cancelled, when its select list contained a complex expression with Unicode letters and no column alias.

This issue is resolved by correcting the parsing loop so that the query proceeds to planning instead of hanging.

Apache Jira: IMPALA-14551

CDPD-98992: Impala crashes when aes_decrypt receives invalid input

7.3.2.10000

Previously, the aes_decrypt function crashed the `impalad` process when it received an input string shorter than 16 bytes in AES_128_GCM mode.

This issue is resolved by validating undersized input so that the function returns without crashing.

Apache Jira: IMPALA-14666

CDPD-102537: Catalogd deadlock during startup

A deadlock condition previously occurred in Catalogd during startup when the initial global metadata reset operation took place under a heavy workload. The thread performing the initial global reset periodically released the write lock to allow metadata operations on loaded databases to proceed. If a database was not yet loaded, a concurrent thread acquired a read lock and waited for the database to load. However, the initial global reset thread could not proceed because it required that same write lock, resulting in a deadlock.

This issue is now resolved by reversing the order of operations. The getOrLoadTable functionality now waits for the initial database metadata to load for the specified database before it acquires the versionLock_ read lock.

Apache Jira: IMPALA-14949

CDPD-99049: Spurious table-not-found errors when workload management is first enabled

7.3.2.10000

Previously, when workload management was enabled for the first time, the catalog service logged misleading errors that the sys.impala_query_log and sys.impala_query_live tables were not found, even though the tables were created automatically and the feature worked correctly.

This issue is resolved by checking whether the tables exist instead of resetting metadata, which removes the spurious errors.

Apache Jira: IMPALA-14468

CDPD-99017: Partition pruning fails when an IN list compares columns of mismatched types

7.3.2.10000

Previously, partition pruning failed with an IllegalStateException when a query used an `IN` list that compared a partition column with values of a different type, such as a string against a date.

This issue is resolved by validating the types before comparison so that pruning handles mismatched types correctly.

Apache Jira: IMPALA-14569

CDPD-98946: Impala crashes when viewing the web UI details of a failed query

7.3.2.10000

Previously, opening the query details page in the web UI for a query that had failed and didn't generate summary.

This issue is resolved by checking for the summary data before reading it.

Apache Jira: IMPALA-14791

CDPD-98954: Admission control service crashes when processing compressed requests

7.3.2.10000

Previously, the global admission control service (admissiond) crashed while dequeuing compressed admission requests, because a cached buffer was freed while queued entries still referenced it.

This issue is resolved by keeping the buffer valid until all dependent entries finish using it.

Apache Jira: IMPALA-14771

CDPD-100324: CREATE TABLE AS SELECT fails an internal validation check

7.3.2.10000

Previously, a `CREATE TABLE AS SELECT` statement could fail an internal validation check when selecting from an external table with certain properties across databases.

This issue is resolved by introducing a mechanism that asynchronously cleans up cancelled query entries from the state map.

Apache Jira: IMPALA-14875

CDPD-101262: A pool can exceed its maximum memory when query memory limits are not set

7.3.2.10000

Previously, queries in a resource pool could use more memory than the configured pool Max Memory when the pool had no minimum or maximum query memory limit and no MEM_LIMIT query option was set, because Impala allowed unlimited memory instead of applying the pool maximum.

This issue is resolved by enforcing the pool maximum memory even when the query memory limits are not configured.

Apache Jira: IMPALA-14840

CDPD-97080: Memory leak in the admission control service when queued queries are cancelled

7.3.2.10000

Previously, the global admission control service leaked memory when queries waiting in the admission queue were cancelled, because their entries were not removed from the internal state map.

This issue is resolved by introducing a mechanism that asynchronously cleans up cancelled query entries from the state map.

Apache Jira: IMPALA-14605

CDPD-56422: Queries fail when a column comment exceeds 256 characters

7.3.2.10000

Previously, a query failed when a table column had a comment longer than 256 characters, even though the Hive Metastore stored the longer comment.

This issue is resolved by handling oversized comments so that the query runs normally.

Apache Jira: IMPALA-3451

CDPD-107170: AI endpoint validation could allow URLs that were not from an allowed provider

7.3.2.10000

Previously, Cloudera AI checked allowed AI endpoints by looking for the allowed domain anywhere in the URL. As a result, a URL that contained an allowed domain only in its path or query parameters could pass validation, even when its actual hostname was not from an allowed provider.

This issue is resolved by updating URL validation for AI endpoints in Impala. While Impala already restricted requests to an allow list of providers, the validation now uses a stricter exact hostname parsing rather than a simple substring match. This change reduces edge cases where allowed domain strings hidden in URL paths or parameters might bypass the policy.

CDPD-106242: Server-side request forgery and unauthorized schema access through Avro schema URL

7.3.2.10000

Previously, specifying a URL in the avro.schema.url table property made unrestricted HTTP requests and exposed raw file or HTTP response details in error messages.

This issue is resolved by disabling remote HTTP requests by default for the avro.schema.url table property in Impala. If required, you can enable remote HTTP requests by using the avro_schema_url_remote_http_enabled and avro_schema_url_http_allowed_hosts configuration properties.

CDPD-106242: Uncontrolled remote HTTP requests in Avro schema URLs

7.3.2.10000

Previously, requests using local filesystem schemes, such as local:// or hdfs://, in the avro.schema.url table property were unguarded and allowed unauthorized file reads without proper access controls.

This issue is resolved by adding Ranger read permission checks for local filesystem paths in Impala.

Cloudera Runtime 7.3.2.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Impala issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

CDPD-98207: Impala crashing on the Web UI for failed queries

7.3.2

Previously, Impala crashed when you accessed the query summary or JSON plan through the Web UI for queries that failed before execution. This occurred during scenarios such as a Create Table As Select (CTAS) failure or when admission control rejected a query.

This issue is addressed by ensuring the system correctly handles missing execution summaries. This issue is now fixed.

Apache Jira: IMPALA-14791

CDPD-97786: Excessive partition events during table-level operations

7.3.2

Previously, certain table-level operations, such as dropping incremental statistics or setting/unsetting cached properties, triggered an individual `ALTER PARTITION` event for every partition in a table.

This issue is addressed by implementing bulk updates for partitions.

Apache Jira: IMPALA-13599

CDPD-97187: Deprecation warnings in impala-shell with Python 3.11

7.3.2

Previously, when running `impala-shell` by using Python 3.11 or newer, the command output displayed DeprecationWarning messages related to ssl.PROTOCOL_TLS and ssl.match_hostname(). These warnings were triggered by underlying library dependencies.

This issue is now resolved by updating the handling of SSL protocols and validating logic to be compatible with newer Python versions, which eliminates these warning messages from the shell output.

Apache Jira: IMPALA-12219

CDPD-91994: Stale query IDs in catalog logs

7.3.2

Previously, catalog logs for getPartialCatalogObject certain metadata requests displayed incorrect query IDs.

This issue is addressed by ensuring that each request is associated with its correct query ID. The system now automatically clears the identification after the request finishes to prevent stale information from appearing in later logs.

Apache Jira: IMPALA-14494

CDPD-82673: RSASSA-PSS certificate signature schema is now supported for server certificates

7.3.2

Previously, if you used a certificate with the RSASSA-PSS signature algorithm for kRPC communication, the connection failed.

The fix includes an updated OpenSSL function that correctly identifies the hash algorithm for RSASSA-PSS certificates.

Apache Jira: IMPALA-14038

CDPD-89852: Crash when casting timestamp strings with timezone offsets to DATE

7.3.2

Attempting to cast a timestamp string that included a timezone offset (like "+08:00" in "2025-08-31 06:23:24.9392129 +08:00") to the DATE data type would cause a crash.

This issue is addressed by adding a check to ensure that the timestamp string length does not exceed the maximum length of the default date-time format. Longer strings will now use a lazily-created format, which prevents the crash.

Apache Jira: IMPALA-14383

CDPD-89730: Impala daemon crashed during scans with high logging levels

7.3.2

Previously, the Impala daemon experienced a null pointer dereference in the BaseSequenceScanner component when the logging level was set to 2 or higher, leading to crashes in release builds.

This issue is resolved by correcting the pointer handling in the sequence scanner to ensure safe memory access when high-level logging is active.

Apache Jira: IMPALA-14382

CDPD-89346: Enhanced join strategy selection for large clusters

7.3.2

The query planner's cost model for broadcast joins can be skewed by the number of nodes in a cluster. This lead to suboptimal join strategy choices, especially in large clusters with skewed data where a partitioned join was chosen over a more efficient broadcast join.

This issue is now resolved by introducing the broadcast_cost_scale_factor query option as an additional tuning option besides query hint to override query planner decision.

Apache Jira: IMPALA-14263

CDPD-89132: Tables incorrectly dropped by stale HMS events after global metadata invalidation

7.3.2

Previously, a stale event such as DropTable or AlterTableRename post global `INVALIDATE METADATA` command could cause tables to be unintentionally dropped

This issue is resolved by tracking the createEventId as the current HMS event ID for all tables during a global reset.

Apache Jira: IMPALA-14330

CDPD-79111: Authentication failure in impala-shell with 76 character LDAP passwords

7.3.2

Previously, when you used `impala-shell` with the HS2-HTTP protocol and a 76 character LDAP password, the connection failed with a value error.

This issue is resolved by an updated encoding method that handles long password strings without inserting line breaks, ensuring that the authorization header remains valid for the server.

Apache Jira: IMPALA-13746

CDPD-92001: Metadata loading performed sequentially in local catalog mode

7.3.2

Previously, when a query accessed multiple unloaded tables in local catalog mode, Impala triggered metadata loading for those tables sequentially.

This issue is resolved by parallelizing table loading during query compilation. A new startup flag, max_stmt_metadata_loader_threads, is introduced to control the number of threads used for loading metadata, with a default value of 8 threads per query. If only one table requires loading or if the thread pool is unavailable, the system automatically falls back to sequential loading.

IMPALA-14447

CDPD-79241: Incorrect query results for Iceberg V2 tables

7.3.2

Previously, when you ran complex queries involving multiple subqueries on Iceberg V2 tables, the system sometimes returned incorrect results.

This issue is now resolved. The fix includes a new internal mechanism to track and apply count optimizations.

Cookie-Based authentication support for JWT tokens

7.3.2

When JWT tokens are used for authentication, every HTTP request within a session requires token verification. If these tokens have a short lifespan, it can lead to authentication failures and disrupt session continuity.

This issue is now resolved by using authentication cookies, which generally have a longer lifespan (configured through the max_cookie_lifetime_s flagfile option) and can remain valid for the duration of the session. This enables subsequent authentication requests to rely on cookies rather than repeatedly verifying the JWT token.

Apache Jira: IMPALA-13813

CDPD-80798: Stable Catalogd initialization in HA mode

7.3.2

Catalogd initialization previously might timeout to complete in high availability mode. This happened because metadata operations started prematurely, blocking Catalogd from becoming active.

This issue is resolved by ensuring Catalogd determines HA state before starting metadata operations in HA mode. This prevents blocking issues and ensures a stable startup.

Apache Jira: IMPALA-13850

CDPD-83059: Optimized Impala Catalog cache warmup

7.3.2

Impala's Catalogd previously started with an empty cache. This led to slow query startup for important tables and affected high availability failovers.

This issue is resolved by adding new settings to pre-load specific tables into the Catalogd cache in the background. This ensures faster query startup and smoother high availability failovers.

Apache Jira: IMPALA-14074

CDPD-87222: Consistent TRUNCATE operations for external tables

7.3.2

Impala's TRUNCATE operations on external tables previously did not consistently delete files in subdirectories, even when recursive listing was enabled.

This issue is resolved by ensuring Impala uses the HMS API for TRUNCATE operations by default.

Apache Impala: IMPALA-14189, IMPALA-14224

DWX-21855: Impala Executors fail to gracefully shutdown

7.3.2

During graceful shutdown Impala executors wait for running queries to finish up to the graceful shutdown deadline (--shutdown_deadline_s). During graceful shutdown the istio-proxy container on Impala executor pod was getting terminated immediately and as a result the executors were not reachable and were removed from the Impala cluster membership resulting in cancellation of running queries.

This issue is now resolved by making sure istio-proxy container’s lifecycle doesn’t impact executor’s cluster membership.

IMPALA-14263: Enhanced join strategy for large clusters

7.3.2

The query planner's cost model for broadcast joins can be skewed by the number of nodes in a cluster. This could lead to suboptimal join strategy choices, especially in large clusters with skewed data where a partitioned join was chosen over a more efficient broadcast join.

This issue is now resolved by introducing the broadcast_cost_scale_factor query option as an additional tuning option besides query hint to override query planner decision. To set it cluster-wide for all queries, add the following key-value to the default_query_options startup option:

```
broadcast_cost_scale_factor=<less than 1.0>
```

Apache Jira: IMPALA-14263

IMPALA-11402: Fetching metadata for tables with huge numbers of files no longer fails with OutOfMemoryError

7.3.2

Previously, when Impala Coordinator tried to fetch file metadata for extremely large tables (those with millions of files or partitions), the Impala Catalog service would attempt to return all the file details at once. This often exceeded the Java memory limits, causing the service to crash with an OutOfMemoryError.

This issue is addressed by configuring the Catalog service to limit the number of file descriptors included in a single getPartialCatalogObject response. A new configuration flag, catalog_part ial_fetch_max_files, is introduced to define the maximum number of file descriptors allowed per response (with a default of 1,000,000 files).

If a request exceeds this limit, the Catalog service will truncate the response and return metadata for only a subset of the requested partitions. The coordinator is now designed to detect this truncated response and automatically send new batch requests to fetch the remaining partitions until all required metadata is retrieved. This change ensures that the coordinator can successfully fetch and process the metadata for extremely large tables without crashing due to memory limits.

Apache Jira: IMPALA-11402

CDPD-77261: Impala can now read Parquet integer data as DECIMAL after schema changes

7.3.2

Previously, if you changed a column type from an integer (INT or BIGINT) to a DECIMAL using `ALTER TABLE`, Impala could fail to read the original Parquet data files. This happened because the files lacked the specific metadata (logical types) Impala expected for decimals, resulting in an error.

Impala is now more flexible when reading Parquet files following schema evolution. If Impala encounters an integer type but the schema expects a DECIMAL, it automatically assumes a suitable decimal precision and scale, allowing you to successfully query the updated table:

- INT32 is read as DECIMAL(9, 0).
- INT64 is read as DECIMAL(18, 0).

This change supports common schema evolution practices by allowing you to update column types without manually rewriting old data files.

Apache Jira: IMPALA-13625

IMPALA-12927: Impala can now correctly read BINARY columns in JSON tables

7.3.2

Previously, Impala couldn't correctly read BINARY columns in JSON tables, often resulting in errors or incorrect data. This happened because Impala assumed the data was always Base64 encoded, which wasn't true for files written by older Hive versions.

Impala now supports a new table property, 'json.binary.format' (BASE64 or RAWSTRING), and a query option, JSON_BINARY_FORMAT, to explicitly define the binary encoding. This ensures Impala reads the data correctly. If no format is specified, Impala will now return an error instead of risking silent data corruption.

JIRA Issue: IMPALA-12927

Fixed Issues in Apache Iceberg

Review the list of Iceberg issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

CDPD-75676: Iceberg uses Hadoop bulk delete API during compaction

7.3.2.20000

Previously, Iceberg removed files during compaction and related maintenance using individual Hadoop filesystem delete operations. When many files were removed in a single operation, especially on object stores that support bulk delete, this approach slowed compaction and file cleanup.

This issue is now fixed. Iceberg can use the Hadoop filesystem bulk delete API to delete multiple files in fewer requests during compaction. Enable this behavior by setting the Iceberg configuration property iceberg.hadoop.bulk.delete.enabled to true. If bulk delete is enabled but the Hadoop bulk delete API is not available in your deployment, the operation fails with an IllegalStateException.

Apache Jira: HADOOP-18679

CDPD-102848: Iceberg time-travel queries fail with UnsupportedOperationException in Hive after cluster upgrade

7.3.2.20000

Previously, after you upgraded a cluster from Cloudera Runtime 7.3.1.700 to 7.3.2.10000, Hive failed to run Iceberg time-travel queries that use FOR SYSTEM_VERSION AS OF on tables that were migrated from Iceberg format version 2 to format version 3. The query failed during compilation with an UnsupportedOperationException in the Calcite optimizer when column statistics were estimated for missing columns. The issue occurred on the upgrade path and did not occur on new 7.3.2.10000 installations.

This issue is now fixed. Hive correctly loads column statistics for Iceberg tables after upgrade, and time-travel queries return data from older snapshots on migrated tables.

Apache Jira: HIVE-29617

Cloudera Runtime 7.3.2.10000 SP1

CDPD-97613, CDPD-100780: Materialized view rebuild fails for Iceberg tables

7.3.2.10000

Previously, materialized view rebuild failed when the base table was an Iceberg table and the materialized view was implicitly created as an ACID table due to metastore.create.as.acid=true resulting in an error. During rebuild, no transaction started but an entry was created in the MATE RIALIZATION_REBUILD_LOCKS table and not cleared, causing subsequent attempts to fail due to an existing lock entry.

This issues is now fixed by setting the metastore.create.as.acid property value to false.

This issue also occurs when the materialized view is an Iceberg table and is tracked in upstream documentation.

Apache Jira: HIVE-29436

CDPD-104751, CDPD-105219: OPTIMIZE TABLE fails on Iceberg V3 tables migrated from V2

7.3.2.10000

Previously, running OPTIMIZE TABLE on Iceberg V3 tables that were migrated from V2 failed with an IllegalStateException in ColumnLineageGraph during query planning. The failure occurred during column lineage processing for hidden columns on migrated V3 tables, which blocked the optimization cycle needed to clear legacy V2 position delete files on newly promoted V3 tables.

This issue is now fixed. Impala excludes OPTIMIZE TABLE from the column lineage graph. OPTIMIZE TABLE is a maintenance operation that rewrites an Iceberg table to a compact form without changing its content.

Apache Jira: IMPALA-15072

CDPD-95072, CDPD-99123: Iceberg metadata scan scheduled to a dedicated executor causes a crash

7.3.2.10000

Previously, on Impala clusters with dedicated coordinators and executors, certain queries on Iceberg metadata tables could schedule an Iceberg metadata scan fragment to a dedicated executor instead of the coordinator. This caused the executor to crash. The issue could occur when a query plan included a UNION over Iceberg metadata scans, even though an earlier fix required metadata scan fragments to run on the coordinator.

This issue is now fixed. Impala now schedules Iceberg metadata scan fragments, including those created by UNION plans, on the coordinator.

Apache Jira: IMPALA-14638

CDPD-102486, CDPD-103287: INSERT fails on Iceberg V3 tables with IndexOutOfBoundsException

7.3.2.10000

Previously, running INSERT on Iceberg V3 tables from Impala failed during query planning with an IndexOutOfBoundsException such as toIndex = 7. CREATE TABLE and SELECT on the same tables worked as expected. The failure occurred when lineage tracking for queries is enabled.

This issue is now fixed. Impala now plans and executes INSERT on Iceberg V3 tables correctly.

Apache Jira: IMPALA-14952

CDPD-102654, CDPD-105436: Iceberg DELETE hangs when the tuple cache is enabled

7.3.2.10000

Previously, when the Impala tuple cache was enabled and MT_DOP was greater than 2, DELETE statements on Iceberg tables could hang indefinitely and remain in a running state. The issue could occur when you ran multiple DELETE statements on the same table.

This issue is now fixed. Impala now completes Iceberg DELETE statements correctly when the tuple cache is enabled.

Apache Jira: IMPALA-14951

CDPD-105713, CDPD-105716: Iceberg V2 count(*) optimization is incorrectly applied to queries that have constant expressions only and no count(*), causing row loss

7.3.2.10000

Previously, on Iceberg V2 tables that contain delete files, queries that have constant expressions only and no count(*) in the select list (for example, SELECT 1 FROM tbl) returned fewer rows than they should.

This issue is now fixed. Impala no longer applies the Iceberg V2 count(*) optimization to queries that do not include count(*).

Apache Jira: IMPALA-14993

CDPD-102876, CDPD-103045: Time-travel queries can result in UnsupportedOperationException in IcebergScanPlanner

7.3.2.10000

Previously, time-travel queries on Iceberg tables from Impala could result in an UnsupportedOperationException during query planning in IcebergScanPlanner. The failure occurred when a time-travel query had to read files on data nodes that are not used by the data files in the current snapshot.

This issue is now fixed. Impala now successfully plans and executes scans for valid Iceberg snapshots in time-travel queries.

Apache Jira: IMPALA-14970

CDPD-99467: Puffin statistics reading is enabled by default

7.3.2.10000

Previously, the initial support for Puffin statistics disabled reading them by default.

This issue is now fixed. Reading Puffin statistics is enabled by default. Impala queries now take available Puffin statistics into consideration.

Apache Jira: IMPALA-14848

CDPD-72897: Inconsistent handling of NOT NULL fields in Iceberg tables

7.3.2.10000

Previously, since IMPALA-12729, inserting a NULL value into a non-nullable column in an Iceberg table could succeed and Impala wrote the NULL value in the Parquet data file. Query results could be inconsistent.

This issue is now fixed. Impala now does not allow writing NULL values into non-nullable columns.

Apache Jira: IMPALA-13273

CDPD-50404: Optimize MIN(part_col), MAX(part_col), and COUNT(DISTINCT part_col) queries for Iceberg tables

7.3.2.10000

Previously, the OPTIMIZE_PARTITION_KEY_SCANS optimization could optimize MIN(key_column), MAX(key_column), and COUNT(DISTINCT key_column) queries using HMS metadata. For Iceberg tables, partitioning statistics are not stored in HMS and are obtained through the Iceberg API instead, so these queries were not optimized in the same way.

This issue is now fixed. For Iceberg V1 and V2 tables without deletes, Impala now optimizes MIN(key_column), MAX(key_column), and COUNT(DISTINCT key_column) queries using partitioning statistics obtained through the Iceberg API when the partition transform is identity. The optimization partially applies to Iceberg V2 tables with deletes.

Apache Jira: IMPALA-11986

CDPD-101398, CDPD-85719: Query compilation failure in Hive Iceberg with ClassCastException

7.3.2.10000

Previously, queries on Hive Iceberg tables with MV rewrite feature failed during the compilation phase with a ClassCastException error.

This issue occurred due to a bug in the code and was triggered when the main query contained a condition that always resulted to false and MV rewrite was applied.

This issue is now fixed.

Cloudera Lakehouse Optimizer

CDPD-93120, CDPD-91628: In-progress task metadata is cleaned up periodically

7.3.2.10000

Previously, running the task cleanup executor for Cloudera Lakehouse Optimizer cleaned up the tasks that were in progress.

This issue is now fixed. The executor cleans up only the completed and failed tasks.

CDPD-93759: Improper output encoding in API responses

7.3.2.10000

Previously, certain Cloudera Lakehouse Optimizer API endpoints returned content without properly encoding special characters such as <, >, &, and ".

This issue is now fixed.

CDPD-93760: Missing security response headers in API responses

7.3.2.10000

Previously, standard security response headers were not included in Cloudera Lakehouse Optimizer API endpoint responses. The missing headers included Content-Security-Policy, X-Content-Type-Options, and Strict-Transport-Security.

This issue is now fixed.

CDPD-99255: Cloudera Lakehouse Optimizer API returns misleading Missing Bearer Token error (401)

7.3.2.10000

Previously, when you accessed the Cloudera Lakehouse Optimizer REST API through the Knox gateway with a valid bearer token but was not a member of the required authorization groups, the API returned a 401 Unauthorized response with a misleading error message indicating that the bearer token was missing.

This issue is now fixed. The API now returns an appropriate authorization error message.

CDPD-98341: Cloudera Lakehouse Optimizer does not handle deleted Iceberg tables correctly

7.3.2.10000

Previously, when an Iceberg table was dropped, Cloudera Lakehouse Optimizer continued to hold stale policy-table associations in the registry, expose deleted table names through read-only APIs, and dispatch policy evaluation tasks against non-existent tables. This caused misleading metrics, phantom associations in the UI and errors from in-flight tasks targeting dropped tables.

This issue is now fixed.

Cloudera Runtime 7.3.2.100 CHF1

CDPD-84534: New configuration to disable block location retrieval

7.3.2.100

Impala is designed to optimize reads by retrieving file block locations from storage systems like HDFS and Ozone, and scheduling read operations on the same nodes where the data resides. In on-premise environments, Impala nodes are typically colocated with storage nodes. However, in scenarios where they are not colocated, retrieving the block location information is a resource-intensive operation.

This issue is now fixed by introducing new Hadoop configuration properties that allow you to disable the retrieval of block location information during table loading. Disabling the retrieval of block location information can be configured in the following ways:

- Globally, using the Hadoop Safety Valve for core-site.xml advanced configuration snippet.

```
'impala.preload-block-locations-for-scheduling': 'false'
```

- By using a Filesystem scheme to disable it for a specific type of storage.

```
'impala.preload-block-locations-for-scheduling.scheme.hdfs':
 'false'
```

- By authority, to disable block location retrieval for a specific endpoint when multiple storage systems are configured with the same scheme.

```
'impala.preload-block-locations-for-
scheduling.authority.mycluster': 'false'
```

This disables block location retrieval only for URIs as hdfs://mycluster/warehouse/tablespace/....

Apache Jira: IMPALA-14138

CDPD-99853, CDPD-99080: Incorrect results when querying multiple Iceberg snapshots

7.3.2.100

Previously, queries on Iceberg tables that referenced multiple snapshots (branch, tag, or current) returned incorrect results because the query planner read data from the same snapshot across all the table references.

This issue is now fixed. The query planner now correctly resolves snapshots for each table reference, ensuring accurate query results.

Apache Jira: HIVE-29509

Cloudera Lakehouse Optimizer

CDPD-102237: Registry is not marked dirty after uploading a policy definition

7.3.2.10000

Previously, after you uploaded a policy definition, the registry was not marked dirty. As a result, the scheduler did not pick up new policy definitions until you reschedule explicitly.

This issue is now fixed.

CDPD-103101: Task event log files accumulate in HDFS after log ingestion

7.3.2.100

Previously, the task event log files in HDFS required manual monitoring and cleanup after the log files were ingested.

This issue is now fixed. The task event log files from HDFS are cleaned up automatically after the log files are successfully ingested into the Iceberg table.

CDPD-103026: Cloudera Lakehouse Optimizer task ingestion action fails due to strict timestamp parsing in Spark3

7.3.2.100

Previously, the task ingestion action failed during Spark validation step because of the strict timestamp parsing in the Spark 3 service.

This issue is now fixed.

CDPD-101842, CDPD-101706: Cloudera Lakehouse Optimizer event policies are triggered even after the policies are deleted

7.3.2.100

Previously, when you tried to associate the event-based policy with an individual table, the association failed abruptly if the policy was already associated with the entire namespace. After you deleted the policy, the event-based scheduler retained the triggers and continued to schedule the table maintenance, which was incorrect.

This issue is now fixed.

CDPD-101534, CDPD-101213: List policy API returns unauthorized status for wildcard namespace

7.3.2.100

Previously, the GET /namespaces/{namespace}/policies API returned the 401 Unauthorized error HTTP status code when you entered the asterisk (*) wildcard for the {namespace} parameter.

This issue is now fixed.

CDPD-101284, CDPD-100562: The Edit Association page displays all namespaces as associated to the Cloudera Lakehouse Optimizer policy

7.3.2.100

Previously, when you edited an association for a **Cloudera Lakehouse Optimizer** policy, the UI incorrectly displayed all the namespaces as associated with the policy or incorrectly displayed all the namespaces as selected.

This issue is now fixed.

CDPD-101183, CDPD-100274, CDPD-101104: API-related performance issues impact Cloudera Lakehouse Optimizer UI response time

Previously, API-related performance issues impacted the **View policy** page load and the "Create Policy" process response time in the Lakehouse Optimizer UI.

This issue is now fixed.

CDPD-84718: Cloudera Lakehouse Optimizer uses the latest com.google.guava:guava 28.1-jre library dependency.

7.3.2.100

Cloudera Lakehouse Optimizer uses the latest com.google.guava:guava 28.1-jre library dependency.

This issue is now fixed.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Iceberg issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

CDPD-78686: Iceberg tables created in 7.2.17 are not captured in Atlas when using 7.2.18 or 7.3.1 Atlas server

7.3.2

This issue occurred due to incompatibility between Data Services, Hive, and Impala hooks in 7.2.17 and the Atlas server in 7.2.18 and 7.3.1.

This fix resolves the compatibility issue and Iceberg tables created in 7.2.17 are now correctly captured and displayed in the Atlas UI in later versions.

CDPD-97171: Concurrency issues between compaction and concurrent write operations

7.3.2

This fix resolves an issue where compaction conflicted with concurrent write operations causing data corruption, by improving concurrency handling to ensure stable operations and data consistency.

Apache Jira: HIVE-29437

CDPD-74040: Dropping Iceberg table with complex type containing timestamp fails

7.3.2

This fix resolves an issue where dropping an Iceberg table with complex types such as array, map, or struct containing timestamp fields failed due to unsupported Hive timestamp type handling, by ensuring proper type conversion and successful table deletion.

CDPD-89402: Event processing invalidates Iceberg tables due to reload failures

7.3.2

This fix resolves an issue where CatalogServiceCatalog.reloadTableIfExists() resulted in a Clas sCastException during event processing, invalidating Iceberg tables and triggering full table reloads instead of incremental loading.

Apache Jira: IMPALA-14358

CDPD-72383: COUNT(*) optimization returns incorrect results in UNION queries on Iceberg V2 tables

7.3.2

A faulty COUNT(*) query optimization caused incorrect results in UNION queries on Iceberg V2 tables that have delete files, leading to inconsistent query outputs.

The fix corrects the optimization logic to ensure accurate results across UNION queries, including scenarios involving post data changes.

Apache Jira: IMPALA-13756 IMPALA-13249

CDPD-81076: LEFT ANTI JOIN fails on Iceberg V2 tables with Delete files

7.3.2

Queries using a LEFT ANTI JOIN fail with an AnalysisException if the right-side table is an Iceberg V2 table containing delete files. For example, consider the following query:

```
SELECT * FROM table_a a
LEFT ANTI JOIN iceberg_v2_table b
ON a.id = b.id;
```

The error Illegal column/field reference'b.input_file_name' of semi-/anti-joined table 'b' is displayed because semi-joined tuples need to be explicitly made visible for paths pointing inside them to be resolvable.

The fix updates the IcebergScanPlanner to ensure that the tuple containing the virtual fields is made visible when it is semi-joined.

Apache Jira: IMPALA-13888

CDPD-78427: Enable MERGE statement for Iceberg tables with equality deletes

7.3.2

This patch fixes an issue that caused MERGE statements to fail on Iceberg tables that use equality deletes.

The failure occurred because the delete expression calculation was missing the data sequence number, even though the underlying data description included it. This mismatch caused row evaluation to fail.

The fix ensures the data sequence number is correctly included in the result expressions, allowing MERGE operations to complete successfully on these tables.

Apache Jira: IMPALA-13674

CDPD-77773: Tolerate missing data files during Iceberg table loading

7.3.2

This fix addresses an issue where an Iceberg table would fail to load completely if any of its data files were missing from the file system. This TableLoadingException left the table in an incomplete state, blocking all operations on it.

Impala now tolerates missing data files during the table loading process. An exception will only be executed if a query subsequently attempts to read one of the specific files that is missing.

This change allows other operations that do not depend on the missing data—such as ROLLBACK, DROP PARTITION, or SELECT statements on valid partitions—to execute successfully.

Apache Jira: IMPALA-13654

CDPD-78508: Skip reloading Iceberg tables when metadata JSON file is the same

7.3.2

This patch optimizes metadata handling for Iceberg tables, particularly those that are updated frequently.

Previously, if an event processor was lagging, Impala might receive numerous update events for the same table (for example, 100 events). Impala would attempt to reload the table 100 times, even if the table's state was already up-to-date after processing the first event.

With this fix, Impala now compares the path of the incoming metadata JSON file with the one that is currently loaded. If the metadata file location is the same, Impala skips the reload, correctly assuming the table is already unchanged. This significantly reduces unnecessary metadata processing.

Apache Jira: IMPALA-13718

CDPD-82415: TABLESAMPLE clause of the COMPUTE STATS statement has no effect on Iceberg tables

7.3.2

This fix resolves a regression introduced by IMPALA-13737. For example, the following query scans the entire Iceberg table to calculate statistics, whereas it should ideally use only about 10% of the data.

```
COMPUTE STATS t TABLESAMPLE SYSTEM system(10);
```

This fix introduces proper table sampling logic for Iceberg tables, which can be utilized for COMP UTE STATS. The sampling algorithm previously located in IcebergScanNode.getFilesSample() is now relocated to FeIcebergTable.Utils.getFilesSample().

Apache Jira: IMPALA-14014

CDPD-85228: IllegalStateException with Iceberg table with DELETE

7.3.2

Running a query on an Iceberg table fails with an IllegalStateException error in the following scenario:

- The Iceberg table has delete files for every data file (no data files without delete files) AND
- An anti-join operation is performed on the result of the Iceberg delete operation (IcebergDeleteNode or HashJoinNode)

This fix resolves the issue by setting the TableRefIds of the node corresponding to the Iceberg delete operation (IcebergDeleteNode or HashJoinNode) to only the table reference associated with the data files, excluding the delete files.

Apache Jira: IMPALA-14154

CDPD-87405: Error unnesting arrays in Iceberg tables with DELETE files

7.3.2

The following error occurred when unnesting a nested array (a 2D array) from an Iceberg table. This issue was triggered specifically when the table contained delete files for some, but not all, of its data files.

```
Filtering an unnested collection that comes from a UNION [ALL] is
 not supported yet.
```

Reading an Iceberg table with this mixed data and delete file configuration creates a UNION ALL node in the query execution plan. The system had a check that explicitly blocked any filtering on an unnested array.

This fix relaxes the validation check, allowing the operation to proceed if all UNION operands share the same tuple IDs. This ensures the query can successfully unnest the array.

Apache Jira: IMPALA-14185

Cloudera Lakehouse Optimizer

CDPD-91628, CDPD-98014: In-progress task metadata is cleaned up periodically

Previously, when the task cleanup executor daemon for Cloudera Lakehouse Optimizer ran, it cleaned up the tasks that were in progress.

This issue is now fixed. The daemon cleans up only the completed and failed tasks that are not in progress.

CDPD-88828: Duplicate tasks run on table

Previously, Cloudera Lakehouse Optimizer allowed duplicate tasks to run on a table.

This issue is now fixed.

CDPD-94583: Filter expression does not transform as expected and fails during the validation step

Previously, when an Iceberg table was created using a DDL command with a bucket transformation specification, and you created a Cloudera Lakehouse Optimizer policy on the table to perform the compaction or rewrite data files tasks, the tasks failed during the validation step.

This issue is now fixed.

CDPD-97815: Associations are lost after Cloudera Lakehouse Optimizer restarts

7.3.2.10000

Previously, the following issues occurred when Cloudera Lakehouse Optimizer restarted:

- The associations defined using the Cloudera Lakehouse Optimizer UI were lost because they were not persisted.
- The namespace-level associations for the tables were not initialized in the registry.

These issues are now fixed.

Fixed Issues in Kafka

Fixed issues and resolved maintenance items for Kafka are addressed in Cloudera Runtime 7.3.2 and its associated service packs.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.10000 SP1

OPSAPS-79006: Kafka broker recommission command fails

7.3.2.10000

The Kafka broker Recommission role action failed in Cloudera Runtime 7.3.2 because host IP resolution ran when it was not required for the recommission command. This issue is fixed. You can recommission a previously decommissioned broker after you upgrade to Cloudera Runtime 7.3.2.10000 SP1.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Kafka issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

OPSAPS-71679: Ranger resource-based policy creation attempted on each service restart

7.3.2

On each service restart, Kafka, Kafka Connect, and Schema Registry attempted to create Ranger resource-based policies even when the Ranger policy cache was already present. Resource-based policy creation is now skipped when the policy cache file exists, reducing unnecessary Ranger operations on restart.

OPSAPS-75094: Broker rolling restart checks fail on FIPS-enabled clusters with JDK 11 or higher

7.3.2

On FIPS-enabled clusters with JDK 11 or higher, CSD_JAVA_OPTS was not propagated to the Kafka broker rolling restart check scripts, which caused the checks to fail. The rolling restart check scripts now receive the CSD_JAVA_OPTS value on FIPS-enabled clusters.

DBZ-4990: The Debezium Db2 Source connector does not support schema evolution

7.3.2

This issue is fixed. For more information, see DBZ-4990.

None.

OPSAPS-69317: Kafka Connect Rolling Restart Check fails if SSL Client authentication is required

7.3.2

Rolling restart checks for Kafka Connect failed when mutual TLS was required (ssl.client.auth set to required), even though the service was healthy. This issue is fixed.

Fixed Issues in Kudu

Fixed issues for Kudu are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.10000 SP1

CDPD-105447: Buffer overflow during message discard operation

7.3.2.10000

Previously, discarding subprocess messages larger than 4 KB caused a buffer overflow or corrupted subsequent messages on the communication channel. This issue led to channel desynchronization and potential process failures during message handling.

This issue is now resolved. The message discard process correctly bounds read operations to prevent memory corruption and maintain proper channel synchronization.

CDPD-106207: Missing Ranger client logs

7.3.2.10000

Previously, an outdated dependency in the Ranger subprocess prevented Ranger client logs from being generated.

This issue is now resolved by updating the logging framework dependencies, which ensures that logs are generated correctly.

CDPD-94251: Catchup deadlock during Raft log synchronization

7.3.2.10000

Previously, a Raft leader could enter an infinite retry loop when attempting log position synchronization with a lagging peer node. This loop consumed high CPU resources and prevented master nodes from completing the catchup phase when joining a cluster.

This issue is rnow esolved by improving heartbeat management during the join process and properly resetting state prior to master node additions, maintaining cluster stability and successful node join operations.

Apache Jira: KUDU-3762

Cloudera Runtime 7.3.2.100 CHF 1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Kudu issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

CDPD-97140: Range partition recreation for specific table schemas

7.3.2

Previously, dropping and immediately adding back the same range partition failed for certain table schemas. This issue occurred because the system catalog did not correctly normalize range-specific hash schemas when they matched the table-wide hash schema. This lack of normalization caused inconsistencies in how the partition information was stored and interpreted. This issue is now fixed. Kudu now correctly normalizes and stores range partition specifications, ensuring that partitions can be recreated as expected.

CDPD-95855: TLS truststore initialization with FIPS-compliant crypto providers

7.3.2

Previously, in some environments using FIPS-compliant crypto providers, TLS truststore initialization failed. This occurred because of incompatible handling of empty keystore initialization. This issue is resolved by updating the truststore initialization logic to ensure compatibility with FIPS-compliant providers. This allows the TLS setup to complete successfully.

CDPD-94474: Kudu tablet server crashes on RHEL 9

7.3.2

Previously, Kudu tablet servers crashed on Red Hat Enterprise Linux 9 (RHEL 9) x86_64 platforms when the libgcc package was updated to version 11.5.0-11 or newer. These crashes typically manifested as a Segmentation Fault (SIGSEGV) or an Abort (SIGABRT) error under specific workloads. This issue is now resolved.

Apache Jira: KUDU-3736

CDPD-80637: Kudu binary crashes on older CPU architectures

7.3.2

Previously, Kudu binaries could crash with an "Illegal instruction" error when running on older CPU architectures that did not support newer instruction sets. This occurred because the Rock sDB library, which is linked to Kudu, was built on newer hardware without the portability option enabled, leading to compatibility issues when the binaries were deployed on older machines. This issue is now resolved.

Apache Jira: KUDU-3580

CDPD-80637: NullPointerException in Kudu Java client

7.3.2

Previously, the Kudu Java client could experience a **NullPointerException** within the TableLocatio nsCache.get() method. This occurred because the client failed to verify if a cache entry specifically for Kudu master locations existed before attempting to check its status. If the entry was absent, the system would attempt to call a method on a null object, leading to a client crash.

This issue is now resolved. The TableLocationsCache now includes a proper null check, ensuring that the Kudu Java client handles missing cache entries gracefully without triggering an exception.

Apache Jira: KUDU-3698

CDPD-80637: Support for RSASSA-PSS signed certificates in channel bindings

7.3.2

Previously, the system could not correctly determine the hash algorithm for certificates signed with RSASSA-PSS. Because the hash algorithm in RSASSA-PSS is configurable, the existing logic failed to identify it, preventing successful channel bindings when these certificates were used.

This issue is now resolved. This ensures compatibility with modern security standards and allows for secure channel bindings when using these certificate types.

Apache Jira: KUDU-3663

CDPD-80637: Incorrect Content-Type headers for binary webserver responses

7.3.2

Previously, the Kudu webserver incorrectly identified certain binary responses, specifically the /ipki-ca-cert-der endpoint as text/plain instead of the expected application/octet-stream. This occurred because the webserver's internal logic relied on a boolean "styled" flag, which could only distinguish between styled and unstyled text, effectively ignoring the "binary" mode required for certificate downloads and other raw data. This issue is now resolved.

Apache Jira: KUDU-3543

CDPD-80637: Error handling for 'kudu table copy' tool

7.3.2

Previously, the kudu table copy command-line tool would crash if invalid or unexpected input was provided for the --predicates flag. Instead of providing a helpful error message, the tool would terminate abruptly, making it difficult for users to identify typos or syntax mistakes in their predicate definitions.

This issue is now resolved. The tool now includes improved input validation; it handles invalid predicate strings gracefully by reporting an actionable error message and exiting without a crash.

Apache Jira: KUDU-3623

Fixed Issues in Apache Knox

Fixed issues and resolved maintenance items for Apache Knox are addressed in Cloudera Runtime 7.3.2 and its associated service packs.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.10000 SP1

CDPD-76490: Ranger API bulk resource deletion fails when proxied through Knox

7.3.2.10000

Fixed an issue where Ranger API bulk resource deletion failed when the request was proxied through Knox Gateway. Bulk resource deletion now works when proxied through Knox Gateway.

Apache JIRA: RANGER-5577

CDPD-98063: Knox service discovery fails in FIPS-enabled clusters with JDK 17 on RHEL 9

7.3.2.10000

Fixed an issue where Knox Gateway could not connect to Cloudera Manager during initial service discovery or periodic configuration analysis in FIPS-enabled clusters running JDK 17 on RHEL 9. Knox Gateway now connects to Cloudera Manager and completes service discovery in these environments.

Apache JIRA: KNOX-3263

Cloudera Runtime 7.3.2.100 CHF 1

CDPD-104770: YARN ResourceManager REST API requests proxied through Knox Gateway ignore query parameters

7.3.2.100

Fixed an issue where Knox Gateway did not forward query parameters when proxying YARN ResourceManager REST API requests through the YARNUI topology. For example, a request to filter applications by state returned all applications instead of only the requested state. Knox Gateway now forwards query parameters on YARN REST API requests.

Apache JIRA: KNOX-3046

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Knox issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

CDPD-94038: Uploading large files through HUE UI using Knox Gateway fails

7.3.2

Fixed an issue where uploading large files through the Cloudera Data Explorer (Hue) user interface using the Knox Gateway failed. Large file uploads now complete successfully.

CDPD-90879: Knox service discovery does not detect Cloudera Manager cluster-level restart commands

7.3.2

Knox service discovery now detects cluster-level `RollingRestart` and `RestartWaitingForStalenessSuccess` commands in Cloudera Manager and automatically regenerates topologies to reflect configuration changes.

Apache JIRA: KNOX-3194

CDPD-78656: Health test for Knox fails if gateway.client.auth.needed is set to true

7.3.2

Previously, the health test for Knox Gateway failed if the gateway.client.auth.needed parameter was set to true.

This issue is now resolved. For TLS Mutual Authentication to work, you must exclude the health topology. To do this, go to Cloudera Manager Knox Configuration , locate the Knox Service Advanced Configuration Snippet (Safety Valve) for conf/gateway-site.xml field, and add a new entry with the following parameters:

```
Name = gateway.client.auth.exclude
Value = health
```

For more information on excluding the topology, see the Apache Knox Documentation.

CDPD-77233: Knox Token TTL value of -1 set to never expire

7.3.2

Fixed an issue where the Knox Token API raised an UnknownTokenException error if the lifespan value of Knox Token TTL was set to -1. Tokens with a TTL value of -1 are now correctly handled as having no expiration when verifying tokens for authentication.

Apache JIRA: KNOX-3075

Fixed Issues in Livy

Review the list of Livy issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.2.10000 SP1

CDPD-101584: Livy for Spark 3 trust store password on BCFKS clusters

7.3.2.10000 SP1

Fixed Livy for Spark 3 failures when the JVM default trust store is BCFKS and no trust store password is set. Livy supports livy.ssl.defaultTrustStorePassword, which sets javax.net.ssl.trustStorePassword on the Livy server JVM when needed. Cloudera Manager may set this property on managed clusters.

CDPD-107418: Livy jobs fail connecting to Hive Metastore over SSL on FIPS clusters

7.3.2.10000 SP1

Fixed SSL handshake failures when Livy Spark sessions connect to Hive Metastore on FIPS 140-3 clusters where the JVM trust store uses BCFKS.

CDPD-110467: Livy cluster-mode cloud storage access fails with empty credential provider path

7.3.2.10000 SP1

Fixed failures when Livy submits Spark jobs in cluster mode to read AWS S3, Azure ADLS, or Google Cloud Storage through Knox IDBroker after an empty hadoop.security.credential.provider.path was set in the Spark environment.

Cloudera Runtime 7.3.2.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.2.0

Cloudera Runtime 7.3.2 resolves Livy issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Livy Fixed Issues.

CDPD-79344: Fixed Knox Gateway gzip exception with Livy HA mode

7.3.2.0

Fixed a bug in Knox Gateway's HA mode where gzip-encoded requests caused duplicated requests and unnecessary failovers. Gzip encoding is now handled more robustly, preventing the un-gzipping error and reducing unnecessary failovers.

CDPD-67187: Fixed missing JVM class imports for Spark

7.3.2.0

JVM classes are now imported during session initialization.

CDPD-63573: Zookeeper ACL support for Livy

7.3.2.0

Added Zookeeper ACL support to Livy's High Availability mode.

CDPD-49717: Fixed issue with Livy not starting if HDFS is still in safe mode

7.3.2.0

Added checks for HDFS safe mode status during Livy server restart, with retries controlled by livy .server.hdfs.safe-mode.interval and livy.server.hdfs.safe-mode.max.retry.attempts.

OPSAPS-69748: Update Ozone filesystem jar name in Livy classpath

7.3.2.0

Updated the Livy classpath to reflect the renamed Ozone filesystem jar.

Fixed Issues in Navigator Encrypt

Review the list of Navigator Encrypt issues that are resolved in Cloudera Runtime 7.3.2.0, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no fixed issues in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.2

The fixed issues for Navigator Encrypt in Cloudera Runtime 7.3.2.0 include all cumulative fixes from lower versions, specifically ranging from Cloudera Runtime 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes, see Fixed Issues.

CDPD-89408: NavEncrypt compile issue with kernel 5.14.0-570

7.3.2

Fixed an issue where the NavEncrypt kernel module failed to compile and install on Red Hat with kernel version 5.14.0-570.35.1.el9_6.x86_64 due to incompatible inode operation definitions.

CDPD-87940: Crash in Navigator Encrypt kernel module if --profile-file= option is used with ACLs

7.3.2

When you added ACLs with the `acl --add --rule` command, if you included any ACLs that used the --profile-file= option, the kernel module (and host) crashed during Navigator Encrypt shutdown.

The issue has been fixed now.

CDPD-81638: ztrustee-c-api memory corruption getting uuids with large replies

7.3.2

Fixed an issue in the ztrustee C API where get deposit operations for UUIDs with large responses could cause memory corruption, potentially leading to runtime errors on some platforms.

Fixed Issues in Oozie

Fixed issues and resolved maintenance items for Oozie are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.2.10000 SP1

CDPD-104226: Spark3 actions fail to connect to Hive due to masked trust store passwords

7.3.2.10000

Previously, Oozie provided the hive-site.xml content to Spark3 actions from a file that is directory generated by Cloudera Manager for Oozie servers. This version of the file often contained masked trust store passwords (example, ********), causing Spark3 actions to fail when connecting to Hive enabled with a java.security.UnrecoverableKeyException configuration. This issue is now resolved and Oozie provides the hive-site.xml content to Spark3 actions directly from /etc/hive/conf/ on the Oozie server hosts, ensuring that trust store properties are correctly passed.

If the old behavior is required, you can set oozie.service.SparkConfigurationService.hive2.configurations:*=hive-conf in the Oozie Server Advanced Configuration Snippet (Safety Valve) for oozie-site.xml.

CDPD-104576: Oozie failed to connect to HBase when ZooKeeper was configured for TLS-only connections

7.3.2.10000 and 7.3.2.20000

Previously, Oozie failed to connect to HBase when ZooKeeper was configured for TLS-only connections because it did not pass the necessary trust store properties to the HBase client. This resulted in connection timeouts and failures during HBase delegation token fetching and Spark3 actions. This issue is now fixed and Oozie uses the hbase-site.xml configuration and passes the required HBase-ZooKeeper TLS properties directly to the connection client.

CDPD-101770: Oozie Git action failed in FIPS-140-3 environments due to trust store access issues

7.3.2.10000

Previously, Oozie's Git action validated configured GitHub endpoints using the Java default cacerts trust store. In FIPS-140-3 environments, where every trust store was required to be password-protected, the Git action failed because it could not access the protected cacerts file. This issue is now fixed and allows you to configure a password-protected trust store for Oozie's Git action (such as the cacerts file, Cloudera Manager's Auto TLS trust store, or any other custom trust store) by providing the trust store location, type, and password through a .jceks file. This is also available for use in non-FIPS environments.

CDPD-96982: Oozie created an unnecessary and insecure /services ZNode in ZooKeeper

7.3.2.10000

Previously, Oozie incorrectly created an unused /services ZNode in ZooKeeper with insecure ACL permissions (world:anyone:cdrwa), even when secure ZooKeeper configuration was enabled. This occurred due to a deprecated code path that ignored the intended Oozie namespace. This issue is now fixed and Oozie correctly uses only the /oozie/services ZNode for service discovery, so the redundant /services ZNode is no longer created.

Cloudera Runtime 7.3.2.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Oozie issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

CDPD-93325: Hive2Main.java can print out passwords

7.3.2

Previously, Hive2Main was printing full Beeline command arguments to standard output, which might unintentionally expose LDAP credentials and the truststore password. This issue is now fixed.

CDPD-78506:Inconsistent PATH environment variable evaluation for Shell actions

7.3.2

Previously, the PATH environment variable for Oozie Shell actions was evaluated on the Oozie server JVM. This issue is now fixed. The PATH is now evaluated based on the YARN NodeManager host settings and applies directly to the Launcher AM container during execution.. You can revert to the legacy behavior by setting the oozie.action.shell.setup-path-in-oozie-server parameter to true in the oozie-site.xml file.

Additionally, the new the oozie.launcher.<***ACTION_TYPE***>.action.env.<***VARIABLE_NAME***> parameter is introduced. This allows you to define action-specific environment variables to customize the execution environments for each action type in the Launcher AM.

CDPD-78069: Oozie action configuration fails to apply Java options after CDPD-60551

7.3.2

Previously, specific Java options and Hadoop properties were not correctly applied to the Oozie launcher configuration. This issue is now resolved. The new oozie.service.HadoopAccessorService.global.yarn.java-options-keys parameter is now introduced, allowing you to configure a comma-separated list of Hadoop properties to be passed to the Oozie launcher if they are not already set in Oozie using the oozie.launcher.<hadoop property key> property or within its workflows.

The default properties in this list are yarn.app.mapreduce.am.admin-command-opts and yarn.app.mapreduce.am.command-opts. The oozie.LauncherConfigurationInjector.hadoop.search.properties property is now deprecated.

CDPD-77768: Oozie log parameters are not substituted correctly

7.3.2

Previously, Oozie log parameters, for example, {0}, were not substituted when log messages contained apostrophes. The formatting engine treated the text following a single quote as a literal block, preventing parameter replacement. This issue is now fixed. Oozie now automatically escapes single quotes in log messages, ensuring that all parameters are substituted correctly.

CDPD-76135: Schema check database connection throws PSQLException

7.3.2

The schema check now completes successfully during Oozie startup for secure database connections or when custom connection properties are used, preventing warnings from appearing in the logs.

CDPD-68425: Oozie does not validate the backend database identifiers in a case-insensitive way

7.3.2

Previously, Oozie incorrectly reported missing tables due to issues with table name casing. This issue is now fixed.

Fixed issues in Ozone

Fixed issues and resolved maintenance items for Ozone are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.2.10000 SP1

CDPD-67848: Rocks tools native libraries are incorrectly packaged inside the Java JAR file

7.3.2.10000

Previously, the Rocks tools native libraries were packaged inside the Java JAR file. Because native libraries were platform-specific while JAR files were platform-agnostic, this caused potential incompatibilities when the same JAR was shipped across different operating systems. This issue is now resolved and moving the shared library files outside of the JAR ensures that native libraries load from the correct platform-specific paths.

CDPD-74200: Ozone Recon fails to synchronize data with Ozone Manager (OM) database

7.3.2.10000

Previously, Ozone Recon failed to synchronize its database with the Ozone Manager (OM) database. This resulted in the Recon UI displaying incorrect or missing information for volumes, buckets, and keys. This issue is now resolved and Recon UI displays accurate namespace data.

CDPD-77957: Ozone Manager follower bootstrap failures with snapshots

7.3.2.10000

Previously, a race condition during the Ozone Manager (OM) bootstrap process could corrupt the database on follower nodes when snapshots were in use. This could cause the OM bootstrap to fail and impact cluster stability. This issue is now fixed and introduces a lock to prevent this race condition, ensuring that the OM bootstrapping is reliable and that the database remains consistent.

Apache JIRA: HDDS-12090

CDPD-78023: OzoneKey isFile returns true for directories

7.3.2.10000

Previously, the OzoneKey.isFile() method incorrectly returned true when called on a directory. This issue is now resolved.

Apache JIRA: HDDS-12094

CDPD-82738: Ozone clients incorrectly attempt to use gRPC ports for streaming when streaming is disabled

7.3.2.10000

Previously, when Ozone streaming was disabled, pipelines did not contain DataNode streaming port information. If a client attempted to use streaming for these pipelines, it incorrectly defaulted to using the gRPC port, resulting in a connection failure. This issue is now resolved.

Apache JIRA: HDDS-12992

CDPD-82805: Container replication fails for over-allocated containers due to inaccurate space reservation

7.3.2.10000

Previously, when exporting or importing containers during DataNode replication, Ozone used a conservative estimate (2x the maximum container size) to reserve disk space. In environments with over-allocated containers, the actual container size exceeded this estimate, leading to inaccurate space accounting and failed replications. This issue is now resolved and ensures 2 * actual container size space is reserved to prevent replication failures for over-allocated containers.

Apache JIRA: HDDS-12998

CDPD-87218: Storage space availability checks for container allocation were performed at a suboptimal stage in the pipeline process

7.3.2.10000

Previously, validation checks for sufficient storage space in a pipeline occurred too early in the allocation workflow, which led to inconsistencies during container creation. This issue is now resolved by moving this logic to the allocateContainer function to ensure that space availability is verified at the most appropriate stage of the allocation process, improving the reliability of disk space handling in Ozone DataNodes.

Apache JIRA: HDDS-13338

CDPD-88421: Ozone Recon UI displays blank data during full OM database snapshot bootstrap

7.3.2.10000

Previously, when Ozone Recon performed a full snapshot fetch of the OM database, it would truncate or clean existing task tables before reprocessing. This caused the Recon UI to appear blank until the synchronization was complete, resulting in a temporary loss of data visibility for users. This issue is now resolved by introducing a staging database mechanism that stores processed data in the background and only updates the final database once reprocessing is finished, ensuring continuous data availability in the UI.

Apache JIRA: HDDS-13515

CDPD-88636: DataNode decommissioning fails when other DataNodes are offline during Erasure Coded (EC) replication

7.3.2.10000

Previously, when attempting to decommission a DataNode while other DataNodes holding Erasure Coded (EC) blocks were offline, the process failed. This occurred because the Storage Container Manager (SCM) attempted to use an affinity node that had been removed from the network topology, resulting in a placement error that prevented decommissioning from completing. This issue is now resolved.

Apache JIRA: HDDS-13544

CDPD-89023: Ozone Tiering profile parameters incorrectly accept negative values

7.3.2.10000

Previously, the Ozone Tiering API accepted negative values (such as -1) for the lessThanEqual and keySizeAboveBytes parameters without returning an error. This caused confusion because these values were treated as "not set" rather than being validated. This issue is now resolved, and the API correctly validates these fields, ensuring that only values greater than or equal to 0 are accepted.

CDPD-89072: The `ozone admin container info` command did not validate container IDs before execution

7.3.2.10000

Previously, the `ozone admin container info` command did not validate the syntax of container IDs passed through the command line before attempting to process them. This led to unnecessary processing or delayed error reporting for invalid IDs. This issue is now resolved, and the command now validates all container IDs upfront and fails immediately if any ID with invalid syntax is provided, ensuring faster feedback and more robust execution.

Apache JIRA: HDDS-13592

CDPD-90077: Ozone Tiering workflows fail immediately when the Ozone Manager (OM) Leader node is unavailable

7.3.2.10000

Previously, Ozone Tiering workflows would fail immediately without any retry mechanism if they encountered transient external dependency issues, such as the Ozone Manager (OM) Leader node being unavailable during an election or maintenance. This required manual intervention to resume the workflow and could cause up to a 24-hour delay in data processing. This issue is now resolved, and a unified retry and timeout framework has now been implemented, allowing tasks to automatically retry with exponential backoff when transient errors occur, improving the overall resilience of the tiering service.

CDPD-91077: Ozone Recon fails to synchronize with Ozone Manager due to duplicate metrics source error

7.3.2.10000

Previously, Ozone Recon encountered a MetricsException stating that the metrics source userTabl eCache already existed during the synchronization process with the Ozone Manager (OM). This occurred when Recon attempted to initialize a new RocksDB store for a full snapshot, leading to

a failure in closing the database and subsequent Rocks Database is closed errors during delta updates. This issue is now resolved by ensuring that metrics sources are correctly managed during the re-initialization of metadata tables.

Apache JIRA: HDDS-13710

CDPD-91507: Mis-replication check fails when a maintenance node becomes unreachable

7.3.2.10000

Previously, when a node in maintenance mode became unreachable and was removed from the network topology, Ozone container placement validation failed. This occurred because the system attempted to check the node's containers for mis-replication using the topology from which the node had already been removed. This issue is now resolved by ensuring the node's network location is correctly referenced even if it is no longer present in the active topology.

Apache JIRA: HDDS-13762

CDPD-91511: Duplicate purge requests in Ozone could lead to incorrect bucket quota usage calculations

7.3.2.10000

Previously, the KeyDeletingService and DirectoryDeletingService services in Ozone did not track inflight purge requests in memory. This resulted in duplicate purge requests being submitted, which led to inaccurate bucket quota usage reporting. This issue is now resolved by tracking the last purge transaction information in memory to prevent redundant requests and ensure accurate storage accounting.

Apache JIRA: HDDS-13760

CDPD-91529: Ozone services fail to export traces to Jaeger due to invalid endpoint format

7.3.2.10000

Previously, Ozone services (such as OM, SCM, DN, and S3G) failed to export traces to Jaeger because the OTEL_EXPORTER_OTLP_ENDPOINT variable was not configured as a full URL. This resulted in an IllegalArgumentException stating that the endpoint must start with http:// or https://. This issue is now resolved.

Apache JIRA: HDDS-13741

CDPD-91531: Excessive logging of full message payloads during WriteChunk or PutBlock failures

7.3.2.10000

Previously, when a client failed to send a WriteChunk or PutBlock request, the system logged the entire message payload, including verbose checksum data. This resulted in excessively large and chatty log files. This issue is now resolved.

Apache JIRA: HDDS-13739

CDPD-91594: Snapshot bucket quota usage is not correctly reduced during background purge operations

7.3.2.10000

Previously, when the background KeyDeletingService and DirectoryDeletingService services purged keys or directories, the SnapshotUsedBytes and snapshotUsedNamespace values were not decreased. This resulted in inaccurate bucket quota tracking and reporting when snapshots were in use. This issue is now resolved, and quota usage is correctly updated during purge requests.

Apache JIRA: HDDS-13764

CDPD-92382: Ozone Manager DB Insights Open Keys table displays incorrect replication information for Erasure Coded (EC) keys

7.3.2.10000

Previously, the Open Keys table in Ozone Manager DB Insights displayed confusing or partial information for Erasure Coded (EC) keys because it used separate columns for replication type and

factor. This issue is now resolved and a combined replication column correctly formats and displays EC replication details.

Apache JIRA: HDDS-13835

CDPD-92409: Namespace summary API returns incorrect directory and key counts

7.3.2.10000

Previously, the Namespace summary API (/api/v1/namespace/summary) returned incorrect file sizes and counts for directories. This occurred because metadata propagation logic skipped the immediate parent directory when a new directory containing files was added to the tree, updating only higher-level ancestors. This issue is now resolved.

Apache JIRA: HDDS-13841

CDPD-92656: Ozone DataNode startup failure on co-located hosts due to storage directory locking

7.3.2.10000

Previously, when Ozone SCM or OM HA and DataNodes were co-located on the same host, the DataNode failed to start with an OverlappingFileLockException. This occurred because the DataNode fell back to using the ozone.metadata.dirs path for Ratis storage when dfs.container.ratis.datanode.storage.dir was not explicitly configured, leading to resource contention and directory locking conflicts with other processes. This issue is now resolved by ensuring proper storage directory handling to prevent locking conflicts on colocated hosts.

Apache JIRA: HDDS-13866

CDPD-92660: Recon UI fails to load Disk Usage view with large number of nested directories and keys

7.3.2.10000

Previously, when there were a large number of entries, the Recon UI Disk Usage (DU) page took time to load or populate the tree. This issue is resolved now and the Disk Usage page no longer loads the root path by default. The user also gets a warning in case they try to submit root path DU. The root path DU can be fetched only after explicit submission by the user.

CDPD-92869: Deadlock in Ozone background services causes intermittent timeouts during snapshot and key deletion

7.3.2.10000

Previously, a deadlock occurred in Ozone background services because the bootstrap lock was acquired after a snapshot was already opened. The bootstrap flow held the bootstrap lock while waiting for a snapshot cache lock, while the background service held the snapshot open and waited for the bootstrap lock, resulting in intermittent TimeoutException errors during background tasks such as snapshot and key deletion. This issue is now resolved.

Apache JIRA: HDDS-13889

CDPD-93006: Snapshot read cache lock is not released during runtime exceptions

7.3.2.10000

Previously, if a runtime exception occurred during Ozone Manager (OM) snapshot processing, the snapshot read cache lock could remain held instead of being released. This unreleased lock could stall downstream operations, leading to deadlocks or failures during the OM bootstrap process. This issue has been resolved by ensuring that exception handling routines properly and safely release the read cache lock.

Apache JIRA: HDDS-13904

CDPD-93585: Ozone client retry is inconsistent when a DataNode has some full disks

7.3.2.10000

Previously, when an Ozone DataNode had some full disks while others remained available, the client hung or repeatedly retried for several minutes. This occurred because the client attempted

to access a full disk instead of failing immediately and reallocating a new block, leading to performance delays. This issue is now resolved.

Apache JIRA: HDDS-14040

CDPD-93952: Recon fails to initialize properly after upgrade due to a race condition

7.3.2.10000

Previously, a race condition occurred during the Recon startup sequence when background services and metrics were initialized before the database schema upgrade was finalized. This caused Recon to fail with a SQLSyntaxErrorException because it attempted to query the RECON_TASK_S TATUS table before the required last_task_run_status column was added. This issue is now resolved.

Apache JIRA: HDDS-14046

CDPD-94032: Missing default value for hdds.datanode.kerberos.principal configuration causes authentication failures

7.3.2.10000

Previously, the hdds.datanode.kerberos.principal configuration property in ozone-default.xml had an empty default value. This inconsistency with other Ozone components required manual configuration and caused Kerberos authentication errors, such as Server has invalid Kerberos prin cipal, particularly when executing DiskBalancer commands in secure clusters. This issue is now resolved by providing a standard default value for the DataNode Kerberos principal.

Apache JIRA: HDDS-14064

CDPD-94035: Ozone Manager logs do not identify the command type for large responses

7.3.2.10000

Previously, when the Ozone Manager generated a response exceeding the size threshold, logs indicated only the response size and protocol method (submitRequest). This made it difficult to identify the specific command causing the large response without cross-referencing audit logs. This issue is now resolved.

Apache JIRA: HDDS-14061

CDPD-94262: Recon fails to start during cluster upgrades due to database schema race conditions

7.3.2.10000

Previously, Recon failed to start during cluster upgrades due to race conditions between database schema upgrades and service initialization. Multiple components attempted to access database tables or columns before the upgrade framework successfully added them, resulting in SQLSynta xErrorException errors and startup crashes. This issue is now resolved.

Apache JIRA: HDDS-14079

CDPD-94836: Ozone Tiering workflow tasks fail due to Hive query timeouts

7.3.2.10000

Previously, Ozone Tiering workflow runs (specifically the AggregateAuditRecords task) failed with a timeout error when the external Hive service was slow or unavailable, because the default connection and socket read timeouts were set to 60 seconds. This resulted in inconsistent workflow success rates. This issue is now resolved, and the default Hive client connection and socket read timeout configuration values are increased to improve the reliability of workflow executions that depend on Hive.

CDPD-94970: Redundant object creation in ContainerInfo.fromProtobuf method

7.3.2.10000

Previously, the ContainerInfo.fromProtobuf method in Ozone created a throwaway instance of Cont ainerInfo before returning the final object. This resulted in unnecessary memory allocation and overhead. This issue is resolved now.

Apache JIRA: HDDS-14196

CDPD-95031: Ozone administrator privilege checks are inconsistent across the system

7.3.2.10000

Previously, Ozone administrator checks were inconsistent because some operations directly checked for administrator privileges while others first verified if Access Control Lists (ACLs) were enabled (ozone.acl.enabled). This led to inconsistent behavior compared to HDFS, where permission checks were verified before superuser privileges. This issue has been resolved by ensuring that Ozone now consistently checks if permissions are enabled (ozone.acl.enabled) before validating administrator privileges.

Apache JIRA: HDDS-14207

CDPD-95195: Ozone Manager fails to start due to missing Java annotation dependencies

7.3.2.10000

Previously, the Ozone Manager (OM) failed to start with a java.lang.NoClassDefFoundError: javax/annotation/Priority error when running on Java 11 or higher. This occurred because the RangerOzoneAuthorizer required specific Java annotation and HK2 libraries that were not included in the Ozone classpath. This issue has been resolved by adding the missing dependencies (javax.annotation-api, hk2-api, hk2-utils, and hk2-locator) to the dependency list.

CDPD-95245: Ozone Manager fails to start in non-HA mode with minimal configuration

7.3.2.10000

Previously, the Ozone Manager (OM) failed to start in non-high availability (non-HA) mode when using a minimal configuration. This occurred because the system could not automatically retrieve the SCM address from default port logic, requiring manual configuration of ozone.scm.client.address and ozone.scm.datanode.address. This issue is now resolved.

Apache JIRA: HDDS-14249

CDPD-95648: Ozone Tiering Workflow API incorrectly reports a running status after a task failure

7.3.2.10000

Previously, the Ozone Tiering Workflow API continued to report the workflow state as running: true even if one or more workflow tasks failed. This resulted in a misleading status where the workflow appeared active despite stopping due to an error. This issue is resolved now.

CDPD-95682: Recon displays incorrect pipeline associations for containers

7.3.2.10000

Previously, the Recon UI incorrectly displayed all available pipelines for every container. This issue is resolved now.

Apache JIRA: HDDS-14368

CDPD-95716: File descriptor leak in the Ozone Manager during checkpoint transfers

7.3.2.10000

Previously, a directory stream was not properly closed in the OMDBCheckpointServletInodeBasedXfer.writeDBToArchive() method. This caused file descriptors to accumulate on the Ozone Manager host, potentially leading to performance degradation or service instability over time. This issue has been resolved by ensuring the directory stream is explicitly closed after execution.

Apache JIRA: HDDS-14376

CDPD-95936: Ozone SCM unnecessarily processes block deletion commands for empty files

7.3.2.10000

Previously, Storage Container Manager (SCM) processed block deletion commands for empty files that contained no data blocks, resulting in unnecessary audit log entries and system overhead. This issue is resolved now.

Apache JIRA: HDDS-14418

CDPD-95970: Recon UI displays "N/A" instead of "0 B" for zero-byte values in Summary Cards

7.3.2.10000

Previously, in the Open Keys Summary and Delete Pending Keys Summary cards, fields with a value of 0 (zero bytes or zero count) were displayed as N/A. This made it difficult to differentiate between valid zero values and missing data or error responses. This issue is now resolved, and zero values are explicitly displayed as 0 B or 0.

Apache JIRA: HDDS-14422

CDPD-96003: Redundant transaction deletion and replica-level size updates in Ozone

7.3.2.10000

Previously, transaction deletion logic in Ozone executed twice, causing redundant delete calls to the database. Additionally, transaction size updates were triggered for each replica instead of at the transaction level, resulting in unnecessary processing and potential aggregation errors. This issue is now resolved.

Apache JIRA: HDDS-14429

CDPD-96234: Ozone bucket table cache is not updated with usage statistics upon key commit

7.3.2.10000

Previously, when a key was committed in Ozone, bucket usage statistics (used bytes and namespace) were updated in the backend but were not reflected in the bucket table cache, resulting in stale or inconsistent usage data being reported. This issue is now resolved.

Apache JIRA: HDDS-14467

CDPD-96416: The Containers page incorrectly displays the "No of Keys" label instead of "No of Blocks"

7.3.2.10000

Previously, the **Containers** page in the Ozone UI incorrectly used the label No of Keys to represent block counts. This was inconsistent with the underlying terminology and the legacy UI. This issue is now resolved.

Apache JIRA: HDDS-14492

CDPD-96710: Non-recursive directory deletion fails with the `Directory is not empty` error

7.3.2.10000

Previously, the S3 clients could unexpectedly encounter a `Directory is not empty` error when attempting to recursively delete a directory through the S3 Gateway on Ozone. This was driven by a transient timing race condition that occurred when the Ozone Manager's internal double buffer had not yet flushed child deletion transactions to the database. This issue is fixed now by properly checking key tombstones in cache in checkSubFileExists or checkSubDirectoryExists.

Apache JIRA: HDDS-14600

CDPD-96831: Ozone-tiering build overwrites core Ozone JAR files in the CDH parcel

7.3.2.0 and 7.3.2.10000

Previously, the Ozone-tiering build process bundled all dependencies, including core Ozone and Hadoop JARs, into its output. During CDH parcel assembly, the Ozone-tiering RPM was extracted after the base Ozone RPM, overwriting patch-level Ozone JARs with older or mismatched versions and potentially reverting patch changes. This issue is now resolved.

CDPD-97209: Container Balancer attempts to move unhealthy containers due to inconsistent validation

7.3.2.10000

Previously, ContainerBalancerSelectionCriteria applied a less stringent health check than MoveManager. This allowed unhealthy containers to be selected for balancing, only to be rejected

later by MoveManager's stricter criteria, resulting in wasted cycles and inefficient balancing operations.

This issue has been resolved by unifying the validation logic, ensuring that the same rigorous health checks are applied upfront. Additionally, if a balancing attempt fails, source DataNodes are now correctly returned to the priority queue so they remain eligible for subsequent balancing cycles.

Apache JIRA: HDDS-14614

CDPD-98087: Performance degradation in SCM listContainer API at scale

7.3.2.10000

Previously, the StorageContainerLocationProtocol.listContainer() API experienced significant performance degradation when handling a large number of containers (200K+). Excessive RPC payloads and per-container lock acquisition caused high CPU usage on the Storage Container Manager (SCM) and slowed Recon sync operations. This issue is now resolved.

Apache JIRA: HDDS-14730

CDPD-98452: Ozone Manager followers could serve stale checkpoints during bootstrap requests

7.3.2.10000

Previously, metadata checkpoints could be downloaded from either a leader or follower Ozone Manager (OM) instance. Fetching a checkpoint from a follower during a manual bootstrap or download resulted in stale metadata being served. This issue is now resolved.

Apache JIRA: HDDS-14775

CDPD-98892: Incorrect file size distribution range displayed in Recon UI

7.3.2.10000

Previously, the Recon UI displayed an incorrect file size distribution range (for example, showing 512KB to 1MB for files actually in the GB range) when a large number of buckets existed within a volume. This issue has been resolved, and the Recon UI now correctly calculates and displays the file size distribution range.

Apache JIRA: HDDS-14827

CDPD-99243: In OzoneManagerSyncMetrics class, the incrNumDeltaRequestsFailed method increments the wrong value

7.3.2.10000

Previously, the metrics for Ozone Manager transaction delta request failures in Recon were not accounted for properly because the incrNumDeltaRequestsFailed method incremented the incorrect variable. This issue is resolved now.

Apache JIRA: HDDS-14848

CDPD-99292: Ozone ThrottledAsyncChecker schedules volume checks in a non-thread-safe manner

7.3.2.10000

Previously, a race condition occurred in ThrottledAsyncChecker where multiple threads could simultaneously schedule checks for the same target or volume. This occurred because the internal map tracking checks in progress was updated only after a check was initiated, resulting in redundant checks or undefined behavior when a volume check was overdue or had not yet run. This issue is resolved now.

Apache JIRA: HDDS-14860

CDPD-99365: DataNode health-check latch timeouts cause false-positive volume failures

7.3.2.10000

Previously, the DataNode StorageVolumeChecker used a single latch for all volume health checks. If this latch timed out due to a transient stall, all pending volumes were immediately marked as failed, even if they were healthy, which led to false-positive volume failures. This issue is now

resolved and DataNodes now tolerate per-volume health-check latch timeouts before marking volumes as failed.

Apache JIRA: HDDS-14871

CDPD-99398: DataNodes may crash during volume failure handling due to active RocksDB iterators

7.3.2.10000

Previously, when a storage volume failure occurred, the StorageVolumeChecker closed the RocksDB instance for that volume without notifying background scanner processes, such as BackgroundContainerDataScanner or OnDemandContainerScanner. If these scanners attempted to use an active iterator on the closed database, a JVM crash occurred. This issue is resolved now.

Apache JIRA: HDDS-14800

CDPD-99431: The `ContainerBalancerStatus` command displays incorrect "Max Size to Move per Iteration" value

7.3.2.10000

Previously, the `ContainerBalancerStatus` command often displayed 0GB for the Max Size to Move per Iteration field. This occurred because the value was incorrectly mapped to the Max Datanodes to Involve per Iteration (percentage) field instead of the Max Size to Move field. This issue is resolved now, and the command correctly displays the maximum size to move per iteration.

Apache JIRA: HDDS-14883

CDPD-99831: Ozone DataNode decommission command hangs when monitoring multiple nodes

7.3.2.10000

Previously, the Ozone DataNode decommission command in Cloudera Manager became stuck indefinitely when decommissioning multiple nodes simultaneously. This occurred because the monitoring script used imprecise string-matching logic (grep) that incorrectly handled IP addresses sharing a common prefix (for example, 10.129.39.23 and 10.129.39.233), causing the script to wait for nodes that had already completed decommissioning. This issue is resolved now.

CDPD-100872: Ozone snapshot rename operation is disabled in the CLI

7.3.2.10000

Previously, the Ozone snapshot rename operation was displayed in the Command Line Interface (CLI) but was not fully supported or functional in Cloudera Base on premises, resulting in unexpected behavior when you attempted to rename snapshots. This issue is resolved now, and the snapshot rename functionality has now been disabled from the CLI to prevent unsupported usage.

CDPD-101434: Risk of missing blocks in Ratis pipelines due to initial under-replication

7.3.2.10000

Previously, when data was written through a Ratis pipeline, blocks were occasionally written only to the leader DataNode, despite the system returning a successful write acknowledgment to the client. Consequently, if the leader DataNode failed or went offline before the data was replicated to other nodes, the associated blocks or chunk data went missing, resulting in temporary data unavailability. This vulnerable state persisted until the Storage Container Manager (SCM) background processes detected the under-replicated container and completed the proper replication. This issue is resolved now.

Apache JIRA: HDDS-15052

CDPD-101613: Unhealthy containers with sufficient replicas are not correctly tracked in container reports

7.3.2.10000

Previously, the RatisUnhealthyReplicationCheckHandler failed to increment the UNHEALTHY container count in reports for containers that were sufficiently replicated but contained unhealthy replicas. While under-replicated and over-replicated unhealthy containers were correctly identified,

those with the correct replica count were missed, leading to inaccurate health reporting in Ozone. This issue is resolved now.

Apache JIRA: HDDS-15261

CDPD-101838: Ozone Manager terminates during startup due to race condition

7.3.2.10000

Previously, a race condition between the Ozone Manager (OM) startup flow and the OM snapshot installation flow caused the OM to terminate unexpectedly with an IllegalStateExce ption. This occurred when both processes attempted to start the delegation token secret manager simultaneously. This issue is resolved now.

Apache JIRA: HDDS-15103

CDPD-101970: Ozone client hangs for a long duration during writeStateMachineData failures

7.3.2.10000

Previously, when a failure occurred during the writeStateMachineData operation on an Ozone Datanode (such as a disk space issue or other exceptions), the client would hang for a long duration while the leader node repeatedly retried the operation until the pipeline closed. This resulted in significant delays before the operation was finally marked as a failure. This issue is now resolved by optimizing the failure handling, reducing the client waiting time to a few seconds in most failure scenarios.

Apache Issue: HDDS-15122

CDPD-102272: Ozone DataNode scanner flags healthy containers as UNHEALTHY due to system resource limits

7.3.2.10000

Previously, background container scanners in Ozone incorrectly marked containers as UNHE ALTHY when they encountered system resource issues, such as file-descriptor exhaustion ("Too many open files"), instead of actual data corruption. This led to false positives where healthy containers were flagged due to temporary system limits. This issue is now resolved, and the logic is updated to catch these resource-related exceptions (FileNotFoundException or FileSystemExcept ion), ensuring that containers remain in their current state if the scanner could not perform its check due to system constraints.

Apache JIRA: HDDS-15150

CDPD-102811: Ozone files marked for deletion were not being reclaimed from snapshots

7.3.2.10000

Previously, files marked for deletion in Ozone were often not reclaimed because their contents were perpetually moved to the next active snapshot during snapshot deletion. This caused deleted files to remain in the cluster indefinitely, consuming storage space. This issue is now resolved by enabling snapshot deep cleaning by default, ensuring that deleted keys are correctly moved through the snapshot chain until they are fully purged.

Apache JIRA: HDDS-15187

CDPD-102907: Ozone Tiering service fails to start due to missing configuration classes

7.3.2.10000

Previously, the Ozone Tiering service failed to start with a java.lang.NoClassDefFoundError for org/apache/commons/configuration2/Configuration. This occurred because of a version mismatch in the classpath dependencies. This issue is now resolved.

CDPD-103013: Delay in safemode exit on follower causing intermittent behaviour leading to leader switch test failure

7.3.2.10000

Previously, Ozone Storage Container Manager (SCM) follower nodes could remain stuck in safe mode after a restart. This occurred because the ContainerSafeModeRule would refresh its container list and potentially add newly created containers that DataNodes (DNs) had already registered, leading to a situation where the 99% reporting threshold could not be met. Consequently, safe mode exit was delayed until a new Ratis transaction was applied, causing intermittent failures in leader switch tests and impacting small clusters. This issue has been resolved by ensuring the safe mode rule maintains the container list initialized at startup and only removes deleted containers, correctly updating the threshold and cutoff values.

Apache Issue: HDDS-15238

CDPD-103659: The `ozone admin container list --all` command returns duplicate containers

7.3.2.10000

Previously, the `ozone admin container list --all` could produce duplicate container entries because SCM returned each batch sorted by container lastUsed timestamp, while the CLI paginated using the last container ID in the batch. When the last item in a batch did not have the highest container ID, the next page re-fetched containers that had already been listed. This is now fixed by sorting list results by container ID on the server.

Apache JIRA: HDDS-15305

CDPD-103728: Inconsistent task state transitions in Tiering Workflow API

7.3.2.10000

Previously, the Tiering Workflow API endpoint returned inconsistent task state transitions, causing tasks that were reported as PASSED to later transition to SKIPPED. This issue is resolved now.

CDPD-103890: Ozone Tiering UI remains blank or stuck in a loading state during the key conversion process

7.3.2.10000

Previously, when a data preparation or key conversion task was triggered in the Ozone Tiering UI, the page remained blank or displayed a loading skeleton until the process finished. This occurred because the UI hid existing data during background refreshes and used a long polling interval, leading to page timeouts and a lack of progress feedback. This issue is resolved now.

CDPD-104319: Tiering UI fails to display SKIPPED status and messages for conversion tasks

7.3.2.10000

Previously, when a tiering workflow skipped a task (for example, if all selected keys were already Erasure Coded), the Overview UI did not handle the SKIPPED state. This caused the workflow to appear as Scheduled with empty metrics, obscuring the skip messages provided by the API. This issue is resolved now.

CDPD-104329: Tiering API workflow returns inconsistent task states for Hive Metastore failures

7.3.2.10000

Previously, when the Tiering workflow encountered a Hive Metastore connectivity failure, task statuses were reported inconsistently across different runs and workflow stages. In some instances, tasks were marked as FAILED, while in others they were marked as SKIPPED, despite having the same underlying root cause. This issue is resolved now.

CDPD-104553: Tiering Workflow API incorrectly reports running status after task failure

7.3.2.10000 and 7.3.2.20000

Previously, the Tiering Workflow API continued to report a status of running=true even after a task within the workflow reached a terminal FAILED state. This resulted in an inconsistent reporting state where the workflow appeared active despite the failure. This issue is resolved now.

CDPD-104558: Incorrect task status transition for AggregateAuditRecords

7.3.2.10000

Previously, the AggregateAuditRecords task was incorrectly marked as COMPLETED in the initial execution snapshot even when the Hive Metastore was down, before it eventually transitioned to FAILED. This race condition resulted in inconsistent status reporting during task execution. This issue is resolved now.

CDPD-106341: SCM logs are flooded with WARN messages when Datanode command count information is missing

7.3.2.10000

Previously, Storage Container Manager (SCM) logs were flooded with WARN messages stating "No command count information for datanode... Assuming zero" during block deletion services, even when the condition was expected or non-critical. This resulted in excessive log noise. This issue is now resolved.

Apache JIRA: HDDS-15588

CDPD-107080: Race condition in Ozone block deletion transaction handling

7.3.2.10000

Previously, a race condition occurred between the readOnlyIterator and onFlush operations during Ozone block deletion. The iterator skipped elements already in the deletion queue, but if another thread cleared the queue before the iterator finished, completed transactions were incorrectly reloaded for deletion. This issue is now resolved.

Apache JIRA: HDDS-15665

CDPD-109489: Snapshot defragmentation fails when directly ingesting a single live RocksDB SST file

7.3.2.10000

Previously, when the incremental snapshot defragmentation service identified exactly one candidate SST file for a table, it bypassed the standard rewrite path and attempted to ingest the live RocksDB SST file directly. This caused ingestion failures (such as External file version not found) or installed incorrect data because the direct path bypassed critical bucket-prefix filtering, snapshot comparisons, and tombstone generation. This issue is now resolved.

Apache JIRA: HDDS-15860

CDPD-110020: Ozone EC offline reconstruction could result in data loss due to premature scrubbing of idle RECOVERING containers.

7.3.2.10000

Previously, data loss occurred during EC offline reconstruction when idle RECOVERING containers were scrubbed and auto-recreated as partial CLOSED replicas. This issue is now resolved.

Apache JIRA: HDDS-15791

CDPD-119182: Ozone Recon disk space exhaustion due to leaked checkpoints

7.3.2.10000

Previously, Ozone Recon failed to clean up temporary checkpoints created during task reinitialization. This caused checkpoint files to accumulate in the /var/lib/hadoop-ozone/recon directory, eventually leading to disk space exhaustion and causing Recon to become unhealthy. This issue is now resolved.

CDPD-111367: Recon Overview displays zero counts for volumes, buckets, and keys after upgrade

7.3.2.10000

Previously, when upgrading from 7.1.9 SP2 to 7.3.2, the Recon Overview page displayed zero counts for volumes, buckets, and keys. This occurred because both versions assigned different features to the same layout version number, causing Recon to skip a database schema update that added required columns to the RECON_TASK_STATUS table and silently failing Ozone Manager (OM) data synchronization. This issue is now resolved.

CDPD-111714: Recommissioning a new Ozone Manager node fails during bootstrap

7.3.2.10000

Previously, when adding or recommissioning an Ozone Manager (OM) node, the OM bootstrap process failed if catching up required an OM checkpoint installation from the leader (for example, when leader logs were purged). The checkpoint installation request was rejected because the OM was in the BOOTSTRAPPING state, causing a deadlock during Ratis configuration staging and preventing the new node from joining the cluster. This issue is now resolved.

Apache JIRA: HDDS-16057

Cloudera Runtime 7.3.2.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Ozone issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

CDPD-84457: Recon logs can be flooded by Negative usedBytes WARN messages in large Ozone clusters

7.3.2

Previously, in Ozone Recon, frequent “Negative usedBytes … treating it as 0” messages were logged at the WARN level and could flood Recon logs in large clusters. This issue has been fixed now.

Apache JIRA: HDDS-13220

CDPD-80567: Snapshot garbage collection fails to reclaim storage

7.3.2

Previously, multiple issues prevented the snapshot garbage collection system from identifying and removing deleted data. This issue in now resolved. Improvements to the efficiency and reliability of snapshot garbage collection process ensure that storage is reclaimed in a timely manner, resulting in better overall performance.

Apache JIRA: HDDS-12558

CDPD-84361: KeyDeletingService fails when the key size exceeds Ratis buffer

7.3.2

Previously, when the KeyDeletingService was fetching keys to be deleted based on keyLimitPerTask, the deletion operation failed if the key size exceeded the Ratis buffer limit (default 32 MB). This issue is now fixed. The key deletion operations no longer depend on the Ratis buffer size.

Apache JIRA: HDDS-13213

CDPD-80739: Ozone Recon - Containers page displays incorrect labels for unhealthy containers

7.3.2

Previously, the Ozone Recon UI incorrectly displayed the Number of Keys label instead of the Number of Blocks label for containers in various unhealthy states. This issue is now fixed. The labels now display the correct information.

Apache JIRA: HDDS-12588

CDPD-84620: Ozone Recon returns 500 error ServiceNotReadyException on /keys/open during NSSummary tree rebuild

7.3.2, 7.3.2.10000

Previously, Ozone Recon returned an HTTP 500 error with a ServiceNotReadyException when the /keys/open API was called while the NSSummary tree was being rebuilt or was temporarily inconsistent. This issue is now fixed.

Apache Jira: HDDS-13763

CDPD-87883: The processed_keys_metrics table fails to update when converting deleted keys

7.3.2

Previously, the processed_keys_metrics table failed to record details when the Ozone tiering workflow attempted to convert deleted keys. This occurred because deleted keys lacked required fields, such as replication type or replication factor,. This issue is now fixed, and the processed_keys_metrics table updates correctly.

CDPD-69122: Ozone Manager database checkpoint generation failure

7.3.2

Previously, the Ozone Manager database checkpoint generation failed due to an InterruptedException Unable to process metadata snapshot request during the parallel snapshot operations or cluster restarts. This issue is now fixed.

Apache JIRA: HDDS-10739

CDPD-92017: Lower Ozone versions cannot process ozone.om.group.rights default value

7.3.2

Previously, lower versions of Ozone could not process the ozone.om.group.rights configuration when it was set to READ, LIST. This issue is now fixed by setting the default value to ALL.

CDPD-75981: Default native ACL limits to user and user's primary group

7.3.2

Previously, the default native ACLs for an object, such as volume, bucket, or file, limited to the object owner and owner's primary group. If Ranger was enabled, these ACLs did not take effect, but were saved to KeyInfo regardless. This issue is now fixed.

Apache JIRA: HDDS-11656

CDPD-87831: SCM over-schedules replications to full DataNodes

7.3.2

Previously, Storage Container Manager (SCM) scheduled replication commands to fix under-replication or mis-replication for container moves, decommissioning, and other operations for both Radis and EC containers. SCM checked whether a target DataNode had space equal to twice the container size value before selecting it as the target node for container replication. However, SCM did not account for the pending operation size of the scheduled tasks. Consequently, SCM could over-schedule replications to a target DataNode that did not have enough space. This issue is now fixed.

Apache JIRA: HDDS-13437

CDPD-80178: Missing check for space availability for all DNS while container creation is in pipeline

7.3.2

Previously, if the leader node in the pipeline did not have the capacity to create a new container, it might have returned a container creation failure. If the follower node did not have the capacity to create a new container, it might have failed and repeatedly attempted to find another follower node. This behavior could cause excessive disk space consumption by parallel write blocks through a state machine, resulting in slower write performance and delayed failure responses. This issue is now fixed by checking whether a DataNode has enough space for a new container before allocating one. This improves write performance and reduces container creation failure in scenarios when DataNodes have less than 5GB disk space remaining.

Apache JIRA: HDDS-12468

CDPD-87749: No logs are available about on-demand scan triggering

7.3.2

Previously, no logs or debug information existed to explain why on-demand scans were triggered on the containers. This issue is now fixed, and logs are available specifying the reason for on-demand container scans.

Apache JIRA: HDDS-13423

CDPD-85250: The OzoneTokenIdentifier does not serialize or deserialize correctly

7.3.2

Previously, a null omServiceId was deserialized as an empty string, which caused delegation token cleanup issues in RocksDB. This issue is now fixed w.

Apache JIRA: HDDS-13264

CDPD-82295: AWS S3 DeleteObject failures for FSO bucket keys containing special characters

7.3.2

Previously, AWS S3 DeleteObject could fail for File System Optimized (FSO) bucket keys containing special characters. This issue is now fixed by removing name validation during deletion.

Apache JIRA: HDDS-12911

CDPD-74686: DirectoryDeletion task ignored by Ratis

7.3.2

Previously, directory deletion tasks were ignored by Ratis, leading to repeated deletion retries instead of actual deletion. This issue is now resolved.

Apache JIRA: HDDS-11491

CDPD-74685: Directory deletion fails having millions of directory

7.3.2

Previously, background directory deletion cleanup failed when attempted to delete millions of empty directories because their combined metadata size exceeded the allowed Ratis request size. This issue is now resolved.

Apache JIRA: HDDS-11492

CDPD-87270: Secret key premature expiration and invalidation

7.3.2

Previously, secret keys could expire before the end of a delegation token lifetime causing premature authentication failures. This issue is now fixed. The secret key expiry calculation (hdds.secret.key. expiry.duration) is adjusted to 9 days. This ensures that tokens remain valid for their full configured duration to improve stable authentication.

Apache JIRA: HDDS-13343

CDPD-76523: ozone debug ldb --with-keys key defaults to false instead of true

7.3.2

Previously, the ozone debug ldb --with-keys option defaulted to false when specified without a value and did not print the keys. This issue is now fixed. The option defaults to true when specified without a value and includes keys in the output by default.

Apache JIRA: HDDS-11782

CDPD-84609: The --output-dir option is unavailable for replicas verify command

7.3.2

Previously, the Ozone `debug replicas verify` command did not support the --output-dir option. This issue is now fixed. The --output-dir option is now an optional field for the replicas ver ify command.

Apache JIRA: HDDS-13248

CDPD-76520: DataNode aborts if hdds.datanode.wait.on.all.followers = true

7.3.2

Previously, the DataNode aborted if the hdds.datanode.wait.on.all.followers configuration was set to true. This issue is now fixed.

Apache JIRA: HDDS-11785

CDPD-76501: DataNode Ratis is taking snapshots frequently

7.3.2

Previously, DataNode Ratis was taking snapshots every 5 to 8 seconds causing overhead. This issue is now fixed. The hdds.ratis.snapshot.threshold and hdds.container.ratis.statemachine.max.pending.apply-transactions configuration limits are increased to 100k to avoid taking frequent DataNode Ratis snapshots.

Apache JIRA: HDDS-11773

CDPD-75112: HBase RegionServer crashes due to inconsistency caused by Ozone client failover handling

7.3.2

Previously, the HBase RegionServer crashed due to inconsistencies caused by Ozone client failover handling. This issue is now fixed by making the Ozone Manager client retry idempotent which prevents the client from crashing when encountering inconsistent results.

Apache JIRA: HDDS-11558

CDPD-77938: Local Refresh button for current selected path is missing in the new Ozone Recon UI

7.3.2

Previously, refreshing the Recon UI page reset the current path selection and returned to the root directory, causing loss of context and requiring manual navigation. This issue is now fixed. The new Path Reload button is introduced in the new Recon UI for the **Namespace** page.

Apache JIRA: HDDS-12085

CDPD-77728: Calendar disappears while setting custom date range in the Heatmap page in New Recon UI

7.3.2

Previously, setting the custom date range in the **Heatmap** page of the new Recon UI caused the calendar widget to close unexpectedly. Specifically, clicking the back arrow to navigate to a previous month in the date picker, caused the entire calendar and the drop-down menu to disappear, preventing date selection. This issue is fixed, and the calendar remains visible until a date is selected and confirmed, allowing users to set custom date ranges as intended.

Apache JIRA: HDDS-12044

CDPD-77356: Recon UI displayed identical and duplicate values for Quota Allowed and Quota In Bytes

7.3.2

Previously, in the Ozone Recon UI, the Quota Allowed and Quota In Bytes fields incorrectly displayed the same value. This duplication prevented you from accurately distinguishing between the allocated quota and the actual consumed disk space. This issue is now fixed, and the Recon UI displays the values correctly.

Apache JIRA: HDDS-11987

CDPD-74437: Multiple IOzoneAuthorizer instances might be created during Ratis snapshot installation failures

7.3.2

Previously, if a failure occurred during the installation of a Ratis snapshot after the metadata manager was stopped, multiple instances of the Ozone authorizer could be created and retained

in memory. This led to excessive heap usage and, in some cases, crashes due to long garbage collection pauses, especially in environments with Ranger and Ozone integration. The issue is now fixed, and the old authorizer instances are properly cleaned up, preventing heap exhaustion.

Apache JIRA: HDDS-11472

CDPD-92003: Container Size Count Task showing empty in new Recon UI

7.3.2

Previously, in the Ozone Recon UI, the **Container Size Count Task** page was displayed empty when accessed through the new user interface. This issue is now fixed.

Apache JIRA: HDDS-13821

CDPD-88628: Ozone Recon Overview page does not load until all APIs are loaded

7.3.2

Previously, the Recon **Overview** page waited for all API calls to complete before displaying any results, causing delays and poor responsiveness. This issue is now fixed, and each card on the **Overview** page now loads independently as soon as its corresponding API call resolves. This change improves overall page responsiveness and ensures that API errors only affect the relevant cards, rather than preventing the entire page from loading.

Apache JIRA: HDDS-13542

CDPD-88541: Namespace Usage page becomes blank when Recon DB is missing

7.3.2

Previously, the **Namespace Usage** page could appear blank if the Recon DB was missing during a fresh installation. This issue is now fixed.

Apache JIRA: HDDS-13528

CDPD-88383: Accessing the new Ozone Recon UI through Knox breaks the UI

7.3.2

Previously, accessing the new Ozone Recon UI through a reverse proxy such as Knox caused the UI to break. This issue is now fixed.

Apache JIRA: HDDS-13512

CDPD-56281: Ozone Manager database updates are blocked while Recon is reprocessing all Recon tasks

7.3.2

Previously, when Recon was reprocessing all Recon tasks, Ozone Manager database updates were blocked, which could cause repeated full snapshots and impact performance. This issue is now fixed by allowing Ozone Manager database updates to proceed concurrently with Recon task processing, preventing unnecessary full snapshots and improving system efficiency.

Apache JIRA: HDDS-8633

CDPD-77805: Improper error handling in the NSSummaryTask

7.3.2

Previously, improper error handling in the NSSummaryTask could lead to data inconsistencies in the Ozone Recon. This issue is now fixed, and ensures robust error handling in Ozone Recon.

Apache JIRA: HDDS-12062

CDPD-80826: Ozone Recon fails during the bootstrapping process

7.3.2

Previously, Ozone Recon did not properly handle failures that occurred during the bootstrapping process. This issue is now fixed. If an Ozone Manager (OM) task fails during bootstrapping, Recon now correctly handles and reprocesses the task to ensure a successful start. Additionally, if Recon receives a partial or corrupted OM database tarball, it cleans up the corrupted file and restarts the fetch process from scratch to maintain data consistency and integrity.

Apache JIRA: HDDS-12615

CDPD-76226: The Recon ListKeys API returns an inappropriate HTTP response

7.3.2

Previously, the Recon `ListKeys` API did not return an appropriate HTTP response when an NSSummary rebuild was in progress. This issue is now fixed. The API now returns the 503 (Service Unavailable) HTTP status code to indicate that the service is temporarily unavailable due to the ongoing NSSummary rebuild. This allows clients to properly handle the too busy or try again later scenario.

Apache JIRA: HDDS-11708

CDPD-76248: The default volume choosing policy is not updated correctly in the ozone-default.xml

7.3.2

Previously, the ozone-default.xml file incorrectly listed the RoundRobinVolumeChoosingPolicy as the default volume choosing policy.This policy did not consider available volume space during container creation or replication, which could result in block allocation failures (though retried) or the creation of small containers. This issue is now fixed. The default volume choosing policy is changed to CapacityVolumeChoosingPolicy in the ozone-default.xml file. This ensures that available capacity is now taken into account during container allocation, improving reliability and resource utilization.

Apache JIRA: HDDS-11735

CDPD-73809: Multithreading issues in the ContainerBalancerTask

7.3.2

Previously, the concurrent access to shared data structures in the getCurrentIterationsStatistic method could cause unpredictable errors. This issue is now fixed. Inside the getCurrentIterationsStatistic method, the system now ensures thread safety by synchronizing access to the iterationsStatistic list and using ConcurrentHashMap for concurrent access to maps from findTargetStrategy and findSourceStrategy.

Apache JIRA: HDDS-11386

CDPD-88723: The FSORepairTool fails to distinguish Unreachable and Unreferenced objects

7.3.2

Previously, the FSORepairTool logic to distinguish between Unreachable and Unreferenced objects was incorrect. This issue is now fixed, and the logic is corrected. The unreachable objects are not marked for repair as background cleanup processes will eventually handle them, while objects that are neither reachable nor unreachable are classified as unreferenced and marked for repair.

Apache JIRA: HDDS-13549

CDPD-87575: The `ozone admin container create` command runs forever without kinit

7.3.2

Previously, the `ozone admin container create` command ran indefinitely on secure Ozone clusters with multiple SCM nodes if authentication failed, for example, when kinit was not performed. This issue was specifically observed in SCM HA cluster configurations. This issue is now fixed, and the retry logic is updated to fail fast on authentication exceptions, providing immediate feedback to you instead of hanging.

Apache JIRA: HDDS-13405

CDPD-90362: Container Balancer stop command fails with an error

7.3.2

Previously, the `stopBalancer` command for the Ozone Container Balancer failed with an error if the balancer was already stopped, instead of returning a successful response. This issue is now

fixed. The `stopBalancer` operation is now idempotent and will return success if the balancer is already stopped.

Additionally, a race condition during an SCM leadership change caused the balancer to restart unintentionally due to the persisted state not being updated. This issue is also now resolved. The system correctly persists the stopped state of the balancer, preventing unintended restarts during leadership transitions.

Apache JIRA: HDDS-13694

CDPD-89400: DataNode pipeline closes frequently

7.3.2

Previously, the DataNode (DN) Ratis repeatedly triggered Close Pipeline actions when it identified issues with a pipeline, such as a slow follower, prolonged leader election, or disk failures, even if a close action was already pending in the DN command queue. This could result in excessive close actions being queued on every heartbeat, leading to inefficiency and potential command queue bloat. The issue is now fixed. A check is introduced to ensure that a Close Pipeline action for a specific pipeline is not added to the command queue if one is already pending, preventing redundant triggers and optimizing the signaling mechanism.

Apache JIRA: HDDS-13618

CDPD-80991: Non-administrative users could attempt to perform OM decommission

7.3.2

Previously, non-administrative users could attempt to perform OM decommission, which could lead to unauthorized or unintended changes. This issue is now fixed. Only users with administrative privileges are authorized to perform OM decommission actions, enhancing the security and integrity of cluster management.

Apache JIRA: HDDS-12646

Fixed Issues in Apache Parquet

Review the list of Parquet issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixed issues in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2

CDPD-74102: Migrate Parquet to Jackson 2

7.3.2

This fix migrates Parquet from the deprecated Jackson 1 (org.codehaus.jackson) to Jackson 2 (com.fasterxml.jackson).

CDPD-74103: Migrate dependency from commons-lang (v2) to commons-lang3

7.3.2

This fix replaces all org.apache.commons.lang.* imports with updated org.apache.commons-lang3 .*. This removes dependency on the deprecated commons-lang version 2 library.

CDPD-77502: Upgrade Parquet to version 1.15.1

7.3.2

The Parquet component has been upgraded from version 1.12.3 to version 1.15.1.

Fixed Issues in Phoenix

Fixed issues for Phoenix are addressed in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Phoenix issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

CDPD-76203: Fails to update to Phoenix 5.1.3 with PhoenixNonRetryableRuntimeException

7.3.2

While updating to the Phoenix 5.1.3 version, the `PhoenixNonRetryableRuntimeException` is thrown which states that the `org.apache.phoenix.query.DefaultGuidePostsCacheFactory` class cannot be loaded or instantiated from SquirrelSQL.

This issue is resolved now.

Apache Jira: PHOENIX-7290

Fixed Issues in Ranger

Review the list of Apache Ranger issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

CDPD-121153: Role resource type missing in Ranger Ozone service definition after cluster upgrade

7.3.2.20000

Following a cluster upgrade, existing Ranger Ozone service definitions were not automatically updated with the new schema, causing the role resource type to be missing from the Ozone policy creation interface in the Ranger Admin UI. This issue has been resolved so that the Ozone service definition schema updates properly during upgrade, making the role resource type available as expected.

CDPD-106846: Ranger Kafka plugin fails to send audit logs to Ranger audit server

7.3.2.20000

In previous versions, the Ranger Kafka plugin could fail to send audit events to the Ranger audit server with a MessageBodyWriter not found for media type=application/json exception due to duplicate Jackson and Jersey JARs in the plugin's classpath. This issue has been resolved by removing conflicting broker-delegate dependencies from the Ranger Kafka plugin implementation.

CDPD-92132: Ranger audit logs were not generated for Impala SHOW DATABASES commands

7.3.2.20000

Previously, running the SHOW DATABASES command from the Impala shell or Impala compute did not generate Apache Ranger audit logs. Additionally, users without Ranger permissions to view the default database could still see it in the query output.

This issue has been resolved. Impala now generates Ranger audit events for SHOW DATABASES operations and enforces Ranger authorization policies so unauthorized users cannot see restricted databases (including the default database) in the results.

CDPD-62959: Improved Ranger Admin HTTPS/TLS startup error logging

7.3.2.20000

Previously, when Ranger UI failed to start or listen on configured HTTPS ports due to keystore or key-alias misconfigurations, Ranger Admin generated misleading error messages (such as RangerConfigUtil: webapp URL is not set), obscuring the root cause. This improvement adds explicit startup validation and error logging for HTTPS keystores, passwords, and key alias (ranger.service.https.attrib.keystore.keyalias) mismatches so TLS-related startup failures can be identified quickly.

Cloudera Runtime 7.3.2.10000 SP1

CDPD-25003: Ranger UserSync failed to synchronize Active Directory (AD) groups containing more than 1500 members

7.3.2.10000

This issue is now fixed. A new configuration property ranger.usersync.ldap.largegroupsync has been introduced in ranger-ugsync-site.xml. Set this property to true in ranger-ugsync-site.xml through safety-valve (Cloudera Manager) to enable correct synchronization of large AD groups. The default value is false. After making the change, restart the Ranger UserSync service.

CDPD-79864: Ranger admin access logs supported wrong timezone

7.3.2.10000

Previously, Ranger admin access logs were recorded exclusively in the UTC timezone, regardless of the local environment settings.

This issue is fixed. You can now configure Ranger to display access logs in the local timezone for improved readability and usability. To enable this, add the ranger.accesslog.timezone.enabled property to the Ranger Admin Advanced Configuration Snippet (Safety Valve) for conf/ranger-admin-site.xml in Cloudera Manager > Ranger > Configuration, and set the value to true.

For details, see How to manage logging properties for Ranger Admin.

CDPD-99737: Ranger now validates the doAsUser parameter for impersonation permission

7.3.2.10000

Previously, the RangerJwtAuthHandler accepted the doAsUser value directly from the incoming request and used it to establish the authenticated user identity without performing any authorization check. This could allow unauthorized impersonation.

This issue is now fixed. Ranger validates that the requesting user has impersonation permission before honoring the doAsUser parameter.

CDPD-67103: Rename should not be allowed on a table having row filter or column masking policy

7.3.2.10000

Previously, renaming a table or column with an active row filter or column masking policy was allowed, which caused the existing policies to become invalid for the renamed table. With this fix, ALTER TABLE rename operations are blocked when the source table has an active row filter or column masking policy, ensuring that security policies remain consistently enforced.

Cloudera Runtime 7.3.2.100 CHF 1

CDPD-103750: Ranger now validates impersonation permission for the doAsUser parameter in JWT authentication

7.3.2.100

Previously, RangerJwtAuthHandler accepted the doAsUser value from incoming requests and used it to establish the authenticated user identity without performing any authorization validation. This could allow unauthorized impersonation of other users.

This issue is now fixed. Ranger validates that the requesting user has the required impersonation permission before honoring the doAsUser parameter.

CDPD-103754: Blank resource type created during policy creation setup

7.3.2.100

Previously, when creating a global policy in Ranger Admin, navigating back from the "Configure Access Rules" step to the "Resources" step and then deleting existing resources and adding a new resource type resulted in a blank resource type being created instead of displaying the available resource type options in the dropdown. This prevented users from successfully adding resources and could lead to policy configuration errors.

This issue is now fixed. The resource type dropdown correctly displays the available options when adding a new resource type, even after navigating back and modifying resources in the global policy creation form.

CDPD-100034: Mixed case key names break HDFS encryption after migration from Key Trustee Server to Ranger KMS DB

7.3.2.100

Previously, when migrating encryption keys from Key Trustee Server (KTS) to Ranger KMS DB, Ranger KMS converted all key names and key version names to lowercase. However, HDFS continued to reference the original mixed-case key names and key version names. This mismatch caused HDFS operations such as reading encrypted files to fail with an error like: KeyVersion '<ver sionName>' does not belong to the key '<keyName>'.

This issue is now fixed.

Cloudera Runtime 7.3.2

The fixed issues for Apache Ranger in Cloudera Runtime 7.3.2 include all cumulative fixes from lower versions, specifically ranging from Cloudera Runtime 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes, see Fixed Issues.

CDPD-91064: USER role accounts to access only their own user details

7.3.2

Fixed an issue where Ranger users with only the USER role could retrieve details of admin or keyadmin users with similar usernames via the /service/xusers/users API. The API now restricts USER-role accounts to accessing only their own user details.

CDPD-90358: Hostname verification issue

7.3.2

Fixed an issue in Ranger where hostname verification did not correctly validate the hostname across the certificate chain, improving SSL/TLS security for secured connections.

CDPD-87716: Ranger RAZ on S3 path with encoded characters

7.3.2

Fixed an issue where Ranger RAZ did not correctly evaluate S3 access policies for Iceberg table partitions when the S3 path contained encoded characters (for example, the = in partition directory names encoded by the AWS SDK v2). Because the encoded path was not decoded before policy evaluation, deny policies on those partition directories were not applied as expected. The path is now decoded prior to authorization checks, ensuring that deny policies on Iceberg partition directories are honored correctly.

CDPD-85478: [GDS] Dataset details not visible in Ranger access logs for RMS enforcement

7.3.2

Fixed an issue where, for RMS enforcement cases, Ranger access logs showed only the dataset policy ID but did not include the corresponding GDS dataset details. Audit entries now correctly capture the GDS dataset information for RMS-enforced datasets.

CDPD-82386: Ranger does not clean up stale plugin entries after role deletion or migration

7.3.2

Ranger now provides supported REST APIs to remove stale plugin status entries from Ranger Admin. Previously, plugin records for deleted or migrated service roles were not cleaned up and continued to appear under Audits > Plugin Status, causing confusion about plugin health and forcing risky manual database updates. The fix introduces ranger-admin endpoints to delete obsolete plugin entries either by ID or by attributes (such as service name, plugin host name, and application type), allowing administrators to safely clean up outdated plugin information and keep the Plugin Status view accurate.

CDPD-81148: Ranger enters an infinite loop and generate persistent database errors

7.3.2

Fixed an issue where Ranger could enter an infinite loop and generate persistent database errors when concurrent sessions updated or deleted the same policy, role, or tag. Tasks scheduled to run after the main transaction (for example, policy label updates, role version updates, and tag version updates) are now processed correctly, preventing infinite loops and ensuring that database errors no longer require a Ranger restart to recover.

CDPD-80625: Improved exception handling for RAZ GCP HMAC keys

7.3.2

Resolved an issue in the RAZ GCP HMAC key generation lifecycle by ignoring transient storage exceptions and retrying key retrieval a few times, and improved related logging. This makes HMAC key creation and deletion more reliable for GCP RAZ integrations.

CDPD-79459: RMS full-sync fails with unsupported Hive table location schemes

7.3.2

Previously, Ranger RMS full-sync could fail with an exception when Hive table locations were stored on file systems that RMS does not support (for example, Azure abfss), causing full-sync to stop before completion.

This issue has been fixed by skipping unsupported file schema types when processing table and database metadata during RMS full-sync and delta-sync. Ranger RMS ACL synchronization now continues for supported file schemes, and locations on unsupported file systems are not mapped; instead, an informational message is logged in the RMS server logs.

CDPD-78321: Performance fixes for Ozone plugin

7.3.2

Fixed the following performance issues observed while evaluating policies for multi-level resources:

- RANGER-4893: Improves policy evaluation for multi level resource hierarchies.
- RANGER-4922: Reduces time to find tags associated with multi-level resources.

CDPD-78151: Added configuration to control append mode for HDFS audit writes

7.3.2

A new configuration parameter was added to Ranger to control whether APPEND mode is used when writing audits to HDFS after errors or exceptions. Previously, Ranger would attempt to append to existing HDFS audit files in such cases (falling back to WRITE mode if append was not possible) to avoid generating a large number of small audit files. With this change, administrators can explicitly enable or disable the use of APPEND mode for HDFS audit writes, allowing better control over audit file handling in specific deployment scenarios.

CDPD-77949: CSV injection vulnerability during CSV/Excel export from Ranger Admin

7.3.2

Fixed an issue that could allow CSV injection during CSV and Excel exports from Ranger Admin (CVE-2024-55532). The export of Ranger policies to CSV/Excel has been removed from the Ranger Admin UI, and the affected export APIs have been deprecated.

CDPD-76633: Ranger RMS server throws ConcurrentModificationException

7.3.2

Previously, the Ranger RMS server could enter an unrecoverable error state with a ConcurrentMo dificationException when large service resource mappings were downloaded to the NameNode while Hive metadata changes were being applied. This could lead to follow#on errors when fetching resource#mapping deltas and no clear path to restore the RMS server to a normal state.

The fix ensures that RMS now uses a shallow copy of the service resource mappings before applying deltas, so the mappings are not modified while they are being serialized for download, preventing the ConcurrentModificationException and stabilizing RMS behavior under concurrent load.

CDPD-75532: Remove self node from the resourceTrie only if it has no children, no evaluators and no wildcard-evaluators

7.3.2

When two policies have a common subset of resources and are defined on the same user (or subset of users, through groups or direct users), if one of these policies is modified (on anything: name, resource, user), it is the only one in effect during access evaluation, until a restart of the underlying service.

This issue has been fixed now.

CDPD-74403: Fixed hardcoded parcel path and sql driver in authzmigrator

7.3.2

Fixed an issue where the authzmigrator/authz-export.sh script failed with NoClassDefFoundError : javax/jdo/JDOHelper when a custom parcel directory or non#default database driver was used. The script now uses the Cloudera Manager parcel directory configuration instead of a hard#coded Cloudera parcel path and automatically selects the appropriate JDBC driver based on the database flavor.

CDPD-73935: Fixed an issue where Ranger “federated user” accounts could log in and perform operations

7.3.2

Ranger now validates the federated user type and prevents these external, data#sharing users from logging in to Ranger, ensuring they are used only for metrics, access history, and audit purposes related to data#sharing features.

CDPD-73779: Support a new user type for external users from Data Sharing

7.3.2

Fixed an issue where Ranger did not distinguish users created by the Data Catalog external user registration process from other external users. Ranger now supports a new Federated User type in

the Ranger Admin UI to represent users originating from Data Sharing (Data Catalog), in addition to the existing Internal and External user types.

CDPD-71673: Security Zone policies version increment issue

7.3.2

This fix addresses an issue where updating a resource caused the associated Security Zone policy version to increment by two instead of one.

CDPD-71563: Issue in dedupTag() method

7.3.2

Fixed a logical flaw in Ranger tag de-duplication that could incorrectly remove valid tags during the dedupTag() operation, preventing policy evaluation failures.

CDPD-69631: Disable Atlas service under the policy permission of tag-based policy

7.3.2

Fixed an issue where the Ranger UI incorrectly allowed selecting the Atlas service in tag-based policy permissions, even though tag-based policies are not supported for Atlas. The Atlas service option is now disabled in tag-based policy permissions to prevent misconfiguration.

CDPD-68970: Fixed an inconsistency between the Ranger UI and Policy Creation API

7.3.2

The Policy Creation API now validates input and fails policy creation if the policy contains only empty values or includes [""] or ["null"] in policyItem users, groups, or roles, aligning API behavior with the Ranger UI.

CDPD-68500: Ranger policy create/update accepts duplicate group names

7.3.2

Previously, Ranger allowed duplicate group and role entries to be added to policies when using the public policy API (/service/public/v2/api/policy), even though the Ranger UI blocked such duplicates. This caused policies created or updated via the API to contain repeated group or role names. The API has been updated to remove duplicate user, group, and role entries during policy creation and update, ensuring consistent behavior with the Ranger UI.

CDPD-68297: REST endpoints do not prevent duplicate values for a resource

7.3.2

REST-based policy creation in Ranger now prevents duplicate values for a resource. Previously, REST API calls could create policies with duplicate resource values (for example, a database list like [test_db1, test_db1]), which could result in multiple policies for the same resource. Policy validation has been updated to reject such requests and require all values for a given resource to be unique.

CDPD-67359: Permissions issues while trying to access folders in Hue file browser

7.3.2

Fixed an issue in the Hue File Browser where accessing S3 directories with more than 1000 objects whose names contained the “=” character failed with a “Cannot Access: <s3 Path>” error. The S3 marker parameter encoding has been corrected so that directories with 1000+ such keys can now be listed successfully through RAZ.

CDPD-67269: Support multiple resource sets in a policy

7.3.2

Improved Ranger policy evaluation to fully support multiple resource sets within a single policy, aligning Cloudera Ranger behavior with the upstream Apache Ranger RANGER-3796 enhancement.

CDPD-62008: Ranger ABAC now supports internal user and group attributes

7.3.2

Resolved an issue where Ranger ABAC policies could not leverage certain internal user and group attributes. ABAC policies can now use the following internal attributes for access control, masking, and row-filtering decisions: syncSource, isInternal, and emailAddress.

CDPD-40734: User allowed to insert data into a hive table when there is a deny policy on a table column

7.3.2

A user is allowed to enter data into a table even if there is a deny policy present on one of the table columns.

This issue has been fixed now.

OPSAPS-75602: Issue with RANGER_C719 CSD becoming stale after upgrading Cloudera Manager

7.3.2

Fixed an issue where the RANGER_C719 CSD could become stale after upgrading Cloudera Manager from 7.13.1.600 with Cloudera 7.1.9 to 7.13.2.0 by fixing the following:

- OPSAPS-73498: Added Cloudera Manager side ranger-trino integration changes.
- OPSAPS-73152: Improved Ranger Admin Diagnostic collection command from Cloudera Manager scripts.

OPSAPS-75556: After upgrade from 7.1.9 to 7.3.2.0 dataset field type is set to boolean in solr managed-schema

7.3.2

Fixed an issue where, after upgrading from Cloudera 7.1.9 to 7.3.2, the datasets field in the ranger_audits Solr collection schema was incorrectly set to the boolean type instead of key_lower_case with multiValued="true". This schema mismatch caused Ranger Admin to fail to load the Access Audit page on upgraded clusters. The upgrade process now updates the ranger_audits Solr schema so that the datasets field is created with the correct type and behaves consistently with fresh 7.3.2 deployments.

OPSAPS-71619: Removed the mandatory validation for ranger.ldap.user.dnpattern

7.3.2

Previously, when LDAP was configured as the external authentication type for Ranger Admin, the ranger.ldap.user.dnpattern parameter was mandatory. If it was not set, the Ranger Admin service failed to start, even though this parameter is rarely required and is ignored when LDAP bind DN/password and user search parameters are configured. This has been fixed by removing the mandatory validation for ranger.ldap.user.dnpattern, so the parameter is now optional and the service can start without requiring a dummy value.

OPSAPS-69156: Fixed an issue with Java add-opens/add-modules/add-exports options

7.3.2

Cloudera Manager components now consistently use the --add-opens=, --add-modules=, and --add-exports= syntax for Java options. This avoids cases where options passed via JAVA_TOO L_OPTIONS could be rejected (for example when using --add-opens or --add-exports without =), improving compatibility across different Java runtimes.

OPSAPS-67197: Ranger RMS server shows as healthy without service being accessible

7.3.2

Previously, Cloudera Manager reported the Ranger RMS server as healthy based only on the RMS process (PID), even when the RMS web service was not fully initialized and the service was inaccessible. The health check logic has been updated to use a Cloudera Manager web alert that verifies the Ranger RMS RMS web endpoint instead of relying solely on the PID. This allows Cloudera Manager to more accurately detect when RMS is not accessible and helps users identify RMS availability issues faster.

Fixed Issues in Ranger KMS

Review the list of Ranger KMS issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.2.10000 SP1

CDPD-97933: Ranger KMS now checks for master key existence on GCP before attempting to create one

7.3.2.10000

Previously, when Ranger KMS was configured with a GCP HSM backend, restarting the Ranger KMS service after revoking the cloudkms.cryptoKeys.create permission would fail with a PERM ISSION_DENIED error. This occurred because Ranger KMS attempted to create the master key on every service restart, regardless of whether the key already existed in GCP. Since the create permission is typically revoked after the initial one-time master key creation, subsequent restarts would fail.

This issue is now fixed. Ranger KMS correctly checks for the existence of the master key in GCP before attempting to create a new one, allowing service restarts to succeed even after the create permission has been revoked.

Cloudera Runtime 7.3.2.100 CHF 1

CDPD-100545: Ranger KMS now checks for master key existence on GCP before attempting to create one

7.3.2.100

Previously, when Ranger KMS was configured with a GCP HSM backend, restarting the Ranger KMS service after revoking the cloudkms.cryptoKeys.create permission would fail with a PERM ISSION_DENIED error. This occurred because Ranger KMS attempted to create the master key on every service restart, regardless of whether the key already existed in GCP. Since the create permission is typically revoked after the initial one-time master key creation, subsequent restarts would fail.

This issue is now fixed. Ranger KMS correctly checks for the existence of the master key in GCP before attempting to create a new one, allowing service restarts to succeed even after the create permission has been revoked.

Cloudera Runtime 7.3.2

The fixed issues for Ranger KMS in Cloudera Runtime 7.3.2 include all cumulative fixes from lower versions, specifically ranging from Cloudera Runtime 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes, see Fixed Issues.

CDPD-74162: Missing HSTS Headers for 404 Not found requests port 9494

7.3.2

Fixed an issue where Ranger KMS HTTPS endpoints on port 9494 did not return HTTP Strict-Transport-Security (HSTS) headers for 404 Not Found responses. Ranger KMS now consistently includes HSTS headers on 404 responses, improving security for clients accessing the KMS service.

CDPD-80814: Failed to migrate the Ranger Kms master key if the masterkey is stored with HDP format

7.3.2

Fixed an issue where migrating the Ranger KMS master key from an HDP-formatted master key to HSM failed with a NullPointerException, preventing import of the master key from the

Ranger database. The migration utility now correctly handles the older HDP master key format and completes the migration successfully.

OPSAPS-72766: Ranger KMS tomcat context update

7.3.2

Updated the default Tomcat context for Ranger KMS from /kms to / by changing the ranger.conte xtName property in ranger-kms-site.xml. This aligns the Ranger KMS context path with Cloudera configuration and simplifies access and integration.

Fixed Issues in Schema Registry

Fixed issues and resolved maintenance items for Schema Registry are addressed in Cloudera Runtime 7.3.2 and its associated service packs.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Schema Registry issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

OPSAPS-71866: Schema Registry fails to start when schema.registry.oauth.clock.skew is set in Cloudera Manager

7.3.2

Schema Registry failed to start when the schema.registry.oauth.clock.skew configuration property was set in Cloudera Manager. This issue is now fixed.

OPSAPS-71679: Ranger resource-based policy creation attempted on each service restart

7.3.2

On each service restart, Kafka, Kafka Connect, and Schema Registry attempted to create Ranger resource-based policies even when the Ranger policy cache was already present. Resource-based policy creation is now skipped when the policy cache file exists, reducing unnecessary Ranger operations on restart.

CDPD-73277: Schema Registry configuration with Oracle TLS fails

7.3.2

The TLS-related Oracle database configuration parameters no longer cause issues during the Schema Registry startup if they are set.

CDPD-94371: Schema Registry does not start with MySQL 8.4 on new installations

7.3.2

Fixed an issue that caused startup failure on new Schema Registry installations when using MySQL 8.4.

Fixed Issues in Apache Solr

Review the list of Solr issues that are resolved in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixed issues in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Solr issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

There are no fixed issues in this release.

Fixed Issues in Spark

Review the list of Spark issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.2.10000 SP1

CDPD-101584: Spark History Server fails on FIPS 140-3 clusters when default cacerts is BCFKS

7.3.2.10000 SP1

Fixed failures when the JVM default trust store is BCFKS and no trust store password is set. Spark 3 adds support for spark.ssl.defaultTrustStorePassword, which sets javax.net.ssl.trustStorePassword on driver and executor JVMs and when Spark History Server starts. Cloudera Manager may set this property on managed clusters.

CDPD-87793: Knox redirection fails for running Spark 3 applications when the Spark UI is disabled

7.3.2.10000 SP1

Fixed Knox redirection to Spark History Server tracking URLs when the Spark UI is disabled and Encrypt all ports defaults apply.

Cloudera Runtime 7.3.2.100 CHF1

There are no fixed issues in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Spark issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Spark Fixed Issues.

CDPD-90482: Jersey Library update

7.3.2.0

Updated the Jersey suite of libraries to versions 2.0 and above to ensure JDK 17 compatibility.

CDPD-82979: Added Cloudera-specific configuration for MapPartitionsRDD default determinacy

7.3.2.0

To avoid failures during error handling, a Cloudera-specific configuration was added to change the default determinacy level for any MapPartitionsRDD.

CDPD-80239: Fixed handling of non-deterministic SQL expressions

7.3.2.0

Fixed a bug related to intermittent data corruption and record loss during task failures. Spark now honors the deterministic=false attribute for non-deterministic data types.

CDPD-67187: Fixed missing JVM class imports for Spark

7.3.2.0

JVM classes are now imported during session initialization.

OPSAPS-74346, OPSAPS-74346, OPSAPS-74316: Enhancing Spark security

7.3.2.0

Changed the generation of the spark.yarn.historyServer.address value to use the HTTPS address when SSL/TLS is enabled. The spark3.network.crypto.enabled new configuration property is now available to enable AES-based encryption.

OPSAPS-72254: UCL | FIPS Failed to upload Spark example jar to HDFS in cluster mode

7.3.2.0

Fixed an issue with deploying the Spark 3 Client Advanced Configuration Snippet (Safety Valve) for spark3-conf/spark-env.sh.

For more information, see *New Cloudera Manager configuration parameter spark_pyspark_executable_path has been added to the Livy for Spark 3* in Behavioral Changes In Cloudera Manager 7.13.2.

OPSAPS-76539: Cloudera Manager UI allowed adding multiple Spark 3 service instances within a single cluster

7.3.2.0

Prevented multiple Spark 3 service instances within a single cluster by implementing a maxInstances flag in Cloudera Manager.

OPSAPS-74281: Disabled the live Spark UI when the ENCRYPT_ALL_PORTS feature flag is enabled

7.3.2.0

Enhanced the security of Spark by implementing new default configuration settings:

1. spark.ui.enabled=false: preventing initiation of an HTTP service that can be accessed from external hosts.
2. spark.io.encryption.keySizeBits=256: the default Spark keySizeBits has been increased from 128 to 256

Fixed Issues in Spark Atlas Connector

Review the list of Spark Atlas Connector issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no fixed issues in this release.

Cloudera Runtime 7.3.2.10000 SP1

CDPD-95322: Missing Atlas lineage for Spark Iceberg tables from MERGE INTO

7.3.2.10000 SP1

Spark SQL MERGE INTO statements on Iceberg tables now send lineage to Apache Atlas.

Cloudera Runtime 7.3.2.100 CHF 1

There are no fixed issues in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Spark Atlas Connector issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Spark Atlas Connector Fixed Issues.

CDPD-95934: Column lineage configuration parameter

7.3.2.0

The Spark configuration parameter atlas.spark.column_lineage.enabled is available to enable or disable column lineage processing. For configuration details, see Spark Connector configuration in Apache Atlas.

Fixed Issues in Apache Sqoop

Review the list of Sqoop issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2

There are no new fixes introduced in this release.

Fixed Issues in Streams Messaging Manager

Fixed issues and resolved maintenance items for Streams Messaging Manager are addressed in Cloudera Runtime 7.3.2 and its associated service packs.

Cloudera Runtime 7.3.2.20000 SP2

CDPD-99346: Newly created partitions do not appear in the Streams Messaging Manager LAG view

7.3.2.20000

Streams Messaging Manager now displays newly created partitions in the LAG view for simple consumer groups after a consumer is assigned to the partitions and commits offsets.

Cloudera Runtime 7.3.2.10000 SP1

CDPD-96648: zstd compression type is missing from the compression.type drop-down menu on the Edit Topic form

7.3.2.10000

When you edit a topic in the Streams Messaging Manager UI, zstd is now available in the comp ression.type drop-down menu on the Edit Topic form. Previously, zstd was not listed even though Kafka supports zstd compression.

CDPD-100017: Streams Messaging Manager configuration with Oracle TLS fails

7.3.2.10000

The TLS-related Oracle database configuration parameters no longer cause issues during the Streams Messaging Manager startup if they are set.

CDPD-99371: Prometheus metrics do not function properly with KRaft mode

7.3.2.10000

If the Kafka cluster monitored by Streams Messaging Manager was running in KRaft mode and Streams Messaging Manager was configured to use Prometheus as its metrics store, alerts related to the active controller did not work as expected. In KRaft mode, the broker_activecontrollerc ount metric was reported by KRaft controllers rather than brokers. This change in metric reporting could cause aggregation issues and lead to alerts triggering indefinitely. Additionally, adding the /api/prometheus-metrics endpoint from KRaft hosts to the Prometheus configuration could cause anomalies in other alerts. This issue is fixed.

If you use Prometheus with KRaft, update your prometheus.yml configuration to scrape KRaft controller metrics endpoints. For configuration details, see *Prometheus properties configuration*.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Streams Messaging Manager issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

OPSAPS-75519: Streams Messaging Manager Rest Admin Server fails to start when the keystore or truststore password contains special characters

7.3.2

The Streams Messaging Manager Rest Admin Server did not handle special characters in keystore or truststore passwords correctly in the generated configuration, which prevented the role from starting. This issue is fixed.

OPSAPS-72814: Streams Messaging Manager used incorrect Ranger policy cache directory

7.3.2

Streams Messaging Manager used an incorrect Ranger policy cache directory setting, which led to permission errors when the Kafka Ranger plugin wrote its policy cache. This issue is fixed.

CDPD-70906: NullPointerException no longer displayed when user is not authorized to access Cruise Control

7.3.2

A message with the text NullPointerException is no longer displayed if the logged-in user is not authorized to access Cruise Control.

CDPD-74620: Data Explorer returns 500 Server Error when switching partitions

7.3.2

In the Streams Messaging Manager Data Explorer, switching between partitions now waits for the new partition's offset information before fetching data. This eliminates the 500 Server Error that occurred when the previous partition's offset was incorrectly used.

CDPD-76375: Topic column name is incorrect in the replications table

7.3.2

Previously, on the **Cluster Replications** page, in the replications table, the topic column name was **Source Topic Name** instead of Target Topic Name.

This issue is now fixed and the table name is the correct **Target Topic Name**.

CDPD-91339: Connector tasks that belong to the same worker cannot be selected individually

7.3.2

Fixed an issue where checking a checkbox in the **Tasks** section on a **Connector Profile** page sometimes selected multiple tasks. Now selecting a single task always selects a single task.

CDPD-62315: Add error boundary to catch runtime errors

7.3.2

An error page fallback is added for unexpected runtime UI errors.

CDPD-45542: Excessive decimal places displayed in metrics

7.3.2

Fixed an issue where task metric percentages in the **Tasks** section of a **Connector Profile** page displayed excessive decimal places, causing UI overflow. Metrics now display with one decimal place.

CDPD-88488: The Streams Messaging Manager server fails to start if its keystore password starts with a special character or includes quotation marks

7.3.2

Fixed an issue where using special characters in keystore passwords made the Streams Messaging Manager fail to start.

Fixed Issues in Streams Replication Manager

Fixed issues and resolved maintenance items for Streams Replication Manager are addressed in Cloudera Runtime 7.3.2 and its associated service packs.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new fixes introduced in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves Streams Replication Manager issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

OPSAPS-73480: Streams Replication Manager upgrade migration incorrectly preserves unprefixed replication properties

7.3.2

During an upgrade, the automatic Streams Replication Manager migration process copies replication properties and removes target-specific prefixes. Previously, if a property existed both with and without a prefix, the migration incorrectly preserved the unprefixed value. For example, if the source contained both custom.myprop=1 and myprop=2, the migration preserved myprop=2. This issue is resolved. The migration now correctly prioritizes and preserves the prefixed value (myprop=1).

OPSAPS-72697: Invalid Streams Replication Manager configuration when Cloudera Manager retried start after a failed attempt

7.3.2

Fixed an issue that could result in configuration errors (invalid configuration files) for Streams Replication Manager if Cloudera Manager automatically tried to start the service again when the first try was unsuccessful. This issue is fixed.

Fixed Issues in YARN and YARN Queue Manager

Fixed issues and resolved maintenance items for YARN and YARN Queue Manager are addressed in Cloudera Runtime 7.3.2 and its associated service packs.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixed issues in this release.

Cloudera Runtime 7.3.2.10000 SP1

COMPX-26640: Added SSL/TLS support for Postgres in QueueManager

7.3.2.10000

Previously, Queue Manager did not support TLS/SSL for PostgreSQL JDBC connections, and was unable to establish secure database connections. This issue is now resolved and this fix automatically appends ?ssl=true&sslmode=require to the JDBC URL when a secured connection is configured, and sets the truststore using the CLIENT_TRUSTSTORE_LOCATION and CLIENT_TRUSTSTORE_PASSWORD environment variables.

COMPX-26371: YARN autoscaling fails to trigger upscaling when pending Application Master (AM) reaches resource limits

7.3.2.10000

The YARN autoscaling logic v1 failed to trigger upscaling when pending Application Masters hit the AM resource limit, even though free cluster capacity existed — causing queues to freeze. A new UpscaleServiceV2 component is now introduced and it accounts for pending applications and AM constraints independently of available cluster resources, enabling proper scaling in AM-constrained scenarios.

COMPX-19423: YARN, MapReduce, and Queue Manager communications were not fully encrypted on the wire.

7.3.2.10000

YARN, MapReduce, and Queue Manager did not enforce TLS encryption across all communication ports, leaving some inter-process and HTTP traffic unencrypted.

This issue is now resolved and all YARN, MapReduce, and Queue Manager communications now support encryption on the wire with a minimum of TLS 1.2 and TLS 1.3 support, along with a new inclusive cipher configuration (ssl.server.include.cipher.list) and a Cloudera Manager parameter to set the supported SSL/TLS versions (hadoop.ssl.enabled.protocols). Cloudera Manager also auto generates certificates for YARN application https endpoints when encrypt-all-ports is enabled.

CDPD-109451: YARN Queue Manager UI failed to load due to false detection of mixed resource allocation mode

7.3.2.10000

Previously,Queue Manager incorrectly detected YARN as being in mixed resource allocation mode after a patch upgrade, causing the Queue Manager UI to fail to load. This issue is now resolved to eliminate the false positive detection, allowing Queue Manager to initialize and load successfully after upgrades.

COMPX-26550: Add serialization fallback logic to the RM's JsonProvider

7.3.2.10000

When a class included in a YARN Resource Manager API response was not registered in either of the Replication Manager's (RM) JSON serialization lists, the RM API failed to serialize the response. This issue is now resolved and this fix prevents serialization failures for unregistered classes.

COMPX-26479: Add RMQueueAclInfo to Resource Manager context

7.3.2.10000

Resolved an issue where the YARN queue access API (`/ws/v1/cluster/queues/{queue}/access`) returned a 503 MessageBodyWriter not found error. The issue was caused by a missing serialization configuration for the RMQueueAclInfo object, which is now corrected to ensure proper JSON serialization of queue ACL responses.

COMPX-25092: YARN web proxy (AmIpFilter) allows TRACE/TRACK requests

7.3.2.10000

Previously, the YARN web proxy (AmIpFilter) did not block HTTP TRACE/TRACK requests, allowing them to bypass Spark UI's security rejections and causing security scanners to flag the vulnerability. This issue is now resolved and updates AmIpFilter to explicitly block TRACE/TRACK methods, returning a 405 Method Not Allowed response.

CDPD-101578: ZooKeeper client TLS truststore password exposed through Resource Manager JVM command-line arguments

7.3.2.10000

The YARN Resource Manager passed the ZooKeeper TLS truststore password through Java command-line options, making it visible thorugh the `ps` commands on the host. This fix secures truststore handling by applying the ZooKeeper truststore password from the environment variable when it is omitted from the JVM command line and initializing JVM SSL defaults from ssl-client.xml.

COMPX-27281: Client JVM SSL fails when JDK_JAVA_OPTIONS omits javax.net.ssl.trustStorePassword

7.3.2.10000

Resolved an issue on FIPS-enabled clusters (7.3.2.10000 and higher versions) where Hadoop and HBase client JVMs failed TLS initialization due to a BCFKS calculation error when javax.net.ssl.trustStore and javax.net.ssl.trustStorePassword were missing. This fix adds ssl-client.xml truststore settings to JVM defaults for HDFS dfs client processes, resolving S3A SSLContext.getDefault() BCFKS failures.

COMPX-24920: YARN Resource Manager fails to start after upgrade due to missing child queue names for root

7.3.2.10000

During YARN startup in a High Availability environment, the YARN Resource Manager intermittently failed to start due to race conditions between two ResourceManagers (RM). One RM started up while another was cleaning up with -format-conf-store, displaying a missing child of the root queue.

This issue is now resolved by backporting the upstream fix (YARN-11924) with retry mechanism to ensure both RMs start with valid configuration.

Apache Jira: YARN-11924

COMPX-24546: Queue Manager lacked FIPS 140-3 cryptographic algorithm validation coverage

7.3.2.10000

Queue Manager did not have test coverage to validate that its cryptographic algorithms comply with FIPS 140-3 standards. CryptoSense/SandboxAQ tests are now added to verify FIPS 140-3 compliance across the primary Queue Manager code paths that leverage cryptography.

CDPD-95076: YARN not compliant with FIPS 140-3 on private cloud and FedRAMP clusters

7.3.2.10000

YARN was not compliant with FIPS 140-3 requirements in 7.3.2 SP1 for private cloud and FedRAMP environments, with issues including non-FIPS-compliant cipher suites, insecure truststore password handling through RM command-line arguments, and missing Cloudera SASL provider configuration.

The fix updates YARN to use FIPS 140-3 compatible cipher suites, secures truststore password handling for the Resource Manager, and integrates CryptoSense-validated cryptographic operations to ensure full FIPS 140-3 compliance.

COMPX-23924: YARN services did not support dual-stack/IPv6 networking

7.3.2.10000

YARN services did not support IPv4/IPv6 preferences, preventing proper connectivity in dual-stack network environments. YARN now starts services with the appropriate java.net.preferIPv4Stack and java.net.preferIPv6Addresses flags based on the IP_VERSION configuration setting, enabling correct resolution to IPv4 and/or IPv6 addresses.

COMPX-21536: YARN Job History Server was not accessible over IPv6

7.3.2.10000

Previously, clients could not connect to the Job History Server or its API endpoints using IPv6 addresses. Dual-stack/IPv6 support is now added to HttpServer2, allowing the Job History Server to resolve and bind to both IPv4 and IPv6 addresses based on the IP_VERSION configuration setting.

CDPD-103660: MapReduce jobs fail with ClassNotFoundException due to missing jts-core in MR-framework

7.3.2.10000

After the Parquet upgrade from 1.15.2 to 1.16.0, MapReduce jobs failed with ClassNotFoundException: org.locationtech.jts.io.ParseException. This issue is now resolved and this fix adds jts-core as a dependency in the parquet-hadoop-bundle, ensuring it is available on the container classpath at runtime.

Cloudera Runtime 7.3.2.100 CHF 1

There are no new fixes in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves YARN and YARN Queue Manager issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

CDPD-49702: Error in NodeManager when executing /var/lib/yarn-ce/bin/container-executor

7.3.2

Previously, a job failure occurred because NodeManager returned a No such file or directory error when attempting to run the /var/lib/yarn-ce/bin/container-executor program. This issue is now

resolved and NodeManager is now marked as unhealthy and shut down if it is unable to run the program.

Apache Jira: YARN-11709

COMPX-13401: CapacityScheduler UI queue filter does not work as expected when submitting applications with leaf queue's name

7.3.2

Previously, when submitting applications to YARN using only a leaf queue name, for example, defa ult or custom, instead of the full queue path, for example, root.default, the ResourceManager (RM) and CapacityScheduler UI inconsistently displayed or filtered applications. This led to confusion, as the same queue could be displayed under different names, and applications were not visible under the expected queue filter in the UI.

This issue is now resolved and RM now returns the full queue path regardless of whether the application was submitted with a leaf queue name or a full queue path.

Apache Jira: YARN-11538

COMPX-14637: Missing permissions on NodeManagers local directories on startup

7.3.2

Previously, the NodeManager created its required local directories on startup with the correct 755 permissions only if they did not already exist.

If an administrator created these directories with incorrect permissions, or if the permissions were altered after the NodeManager started, the NodeManager failed to reset them. This lack of permission enforcement caused container failures. This issue is now resolved.

Apache Jira: YARN-11703

COMPX-17261: NodeManager disk health status oscillation with MB-based limits

7.3.2

Previously, the NodeManager disk health status could oscillate rapidly between good and full states when using the yarn.nodemanager.disk-health-checker.min-free-space-per-disk-mb property because it relied on a single threshold. This issue is now fixed. To prevent this oscillation, the yarn.nodemanager.disk-health-checker.min-free-space-per-disk-watermark-high-mb optional configuration property is now available. This setting allows administrators to specify the minimum free space in megabytes required for a previously bad disk to be marked as good again. If this value is not set or is lower than the YARN.nodemanager.disk-health-checker.min-free-space-per-disk-mb value, it defaults to the same value as the existing minimum free space configuration. This update applies to both YARN.nodemanager.local-dirs and YARN.nodemanager.log-dirs directories, providing more granular control over disk health checks.

Apache Jira: YARN-9914

COMPX-18004: Metrics missing in the RM UI2

7.3.2

The Replication Manager (RM) UI2 did not display metrics with the same granularity as RM UI (UI1). This made analyzing and debugging the scheduler behavior difficult, often requiring the retrieval of information from the UI. This issue is now resolved and the missing metrics are now available in RM UI2.

Apache Jira: YARN-11755

COMPX-18545: Setting maximum-application-lifetime using AQCv2 templates does not apply on the first submitted application

7.3.2

Setting the maximum-application-lifetime property using the AQC v2 templates did not apply to the first submitted application but was applied to the subsequent ones. This issue is now resolved.

Apache Jira: YARN-11708

COMPX-18909: NodeManager marked as unhealthy if an application is terminated

7.3.2

By design, Node Managers are marked unhealthy if an unrecoverable configuration error occurrs, that is, the container-executor script is missing. Previously, a falase positive marking occurred if an application application was terminated just before one of its containers was tried to access the localizer syslog file. This caused an IOException, and the Node Managers was incorrectly marked unhealthy. This issue is now resolved. The error checking is more specific, preventing these false positive unhealthy markings.

Apache Jira: YARN-11753

COMPX-21537: Upgraded the Jersey version

7.3.2

Upgraded the Jersey framework from version 1.19 to 2.46 to fix the CVE-2017-1000028.

COMPX-23191: Null Pointer Exception in Delegation Token Renewer causes all subsequent applications to fail

7.3.2

Previously, any uncaught exception in DelegationTokenRenewer.RenewalTimerTask#run caused all subsequent YARN applications to fail with a java.lang.IllegalStateException: Timer already cancelled exception. This issue is now resolved, and such failures are prevented.

Apache Jira: YARN-11384

COMPX-24259: Incorrect permissions on NodeManager local directories cause container failures

7.3.2

The NodeManager creates the configured local directories with 755 permissions on startup if the directories did not exist. However, if these permissions were changed after startup or if an administrator created the directories with incorrect permissions before starting YARN, the NodeManager did not reset the permissions, resulting in container failures. This issue is now resolved.

Apache Jira: YARN-11703

COMPX-21461: The queuemanager_includedCipherSuites property fails when using comma as a separator

7.3.2

Previously, the queuemanager_includedCipherSuites property in Queue Manager only supported colon (:) as a separator for cipher suites. When comma (,) was introduced as an additional separator, configurations using commas caused failures in property parsing.

This issue is now resolved by updating the property parsing logic to accept both colon and comma as valid separators.

COMPX-22209: Missing centralised Apache HttpComponents libraries

7.3.2

Previoulsy, the cpx component did not use the centralized Apache HttpComponents libraries (http core, httpclient, httpcore5, httpclient5, httpcore5-h2). This issue is now resolved. The component now uses these centralized libraries, to align with with internal standards and incorporate the latest fixes.

COMPX-22213: Missing centralised Bouncy Castle (org.bouncycastle) libraries

7.3.2

Previously, the cpx component did not use the centralized Bouncy Castle org.bouncycastle library versions defined in CDPD (bcprov-jdk18on, bcpkix-jdk18on, and bcutil-jdk18on updated from 1.78

to 1.78.1).This issue is now resolved. The component now uses these centralized libraries to align with internal dependency standards and incorporate the latest security and bug fixes.

COMPX-23423: Apache Commons Lang upgraded to 3.18.0

7.3.2

The Apache Commons Lang package is now upgraded to version 3.18.0 in Queue Manager.

CDPD-91280: Tez is unable to start on 7.3.2.0 FIPS clusters

7.3.2

Previously, Tez startup failed on 7.3.2.0 FIPS clusters, caused by a conflict between Cloudera Manager, which set hadoop.security.secret-manager.key-generator.algorithm to HmacSHA2 56, and Tez's older, hardcoded use of the HmacSHA1 algorithm. This issue is now resolved by upgrading Tez to version 0.10.5 and updating its secret managers to dynamically respect the algorithm configured by Hadoop at runtime. This change prevents the recurring DIGEST-MD5: digest response format violation errors.

Fixed CVE-2025-28924

7.3.2

Few Hadoop API endpoints are now removed from Cloudera Runtime 7.3.2. This change is a result of the fix for CVE-2025-48924 and YARN's migration to common-lang 3. For more information see, https://docs.cloudera.com/cdp-private-cloud-base/7.3.2/private-release-notes/topics/rt-api-compat-changes-hadoop.html

Apache Jira: YARN-10772

Fixed Issues in ZooKeeper

Review the list of ZooKeeper issues that are resolved in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no new fixed issues in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no new fixed issues in this release.

Cloudera Runtime 7.3.2.100 CHF 1

CDPD-102464: ZooKeeper configuration properties logged without redaction

Previously, ZooKeeper configuration properties were recorded in plain text, potentially exposing sensitive data within the logs. This issue is now resolved by implementing a log redaction method to ensure that all sensitive values are appropriately masked during the logging process.

CDPD-102463: Reverse DNS lookup in TLS client and server

Previously, ZooKeeper performed reverse DNS lookups during TLS handshakes in both the client and server, which introduced a high security vulnerability.

This issue is now resolved. Reverse DNS lookup is now disabled in the ZooKeeper TLS client and server.

Apache Jira - ZOOKEEPER-4986

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 resolves ZooKeeper issues and incorporates fixes from the service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all fixes in Cloudera Runtime 7.3.1.x, see Fixed Issues.

CDPD-92113: Zookeeper client uses IPv6 address in case of dual#stack cluster

7.3.2

Previously, in some dual-stack environments, the Zookeeper client failed to connect if IPv6 was not reachable. This issue is now resolved to improve the connection reliability for affected configurations.

CDPD-93023, CDPD-91630: ZooKeeper build broken in JDK17 line

7.3.2

Previously, the ZooKeeper JDK17 build failed due to the missing jline dependency. This issue is now resolved.

CDPD-68499: Upgraded jline version to 3.25.1

ZooKeeper is now updated to use the jline version 3.25.1, which addresses CVE-2023-50572.

OPSAPS-76345: Cloudera Manager could not connect to the ZooKeeper JMX port when JMX TLS was enabled

7.3.2

Previously, an issue triggered by JDK upgrades related to CVE-2026-21925, prevented Cloudera Manager from securely connecting to the ZooKeeper JMX port when JMX TLS was enabled (the default for secure clusters). The problem was caused by a change in the JVM. This issue is now resolved.



Note: Upgrading to a Cloudera Manager version with this fix might lead to temporary staleness.

OPSAPS-75121: TLS Protocol and Cipher Selection support for ZooKeeper

7.3.2

Previously, when two options were removed from Cloudera Manager, related to ZooKeeper TLS settings, testing showed that the TLS version and cipher list settings became inconsistent and insecure. This issue is now resolved and the Zookeeper TLS version and cipher list settings are now removed.

OPSAPS-73956: Support IPv6 access for client

7.3.2

The ZooKeeper startup script, zkserver.sh now uses the configurable variable IP_VERSION instead of a hardcoded value. This change allows the script to dynamically set the appropriate network stack preference, enabling support for IPv4, IPv6, or dual-stack environments as configured. No further manual flag adjustments are needed for IPv6 or dual-stack modes; the system automatically handles the correct settings.

OPSAPS-73624: Data Lake upgrade failure due to ZooKeeper StartPreservingDatastore error

7.3.2

Previously, during the Data Lake (DL) Zero Downtime Upgrade (ZDU) process, the upgrade failed at the ZooKeeper service upgrade phase. This blocked the upgrade and required manual intervention to proceed.

This issue in the ZooKeeper start command is now resolved. The upgrade process now handles ZooKeeper server startups more reliably, preventing the configuration mismatch error and allowing the upgrade to complete without manual steps.

OPSAPS-72239: Failure when installing Cloudera 7.3.1 using Cloudera Manager 7.13.1.0

7.3.2

Previously, when selecting the ZooKeeper gateway role during the initial installation of a Cloudera 7.3.1 cluster using Cloudera Manager 7.13.1.0 caused the wizard to fail with an internal error, preventing the cluster setup.

This issue is now resolved and Cloudera Manager now filters the ZooKeeper gateway role during the initial installation wizard, preventing the error and allowing normal cluster completion.

OPSAPS-63656: Duplicate chart displyed in the ZooKeeper UI

7.3.2

Previously, the Outstanding Requests Across Servers chart was displayed twice in the Cloudera Manager ZooKeeper Server Aggregates UI.

This issue is now resolved and the duplicate ZooKeeper chart is now removed.

OPSAPS-57641: ZooKeeper Authentication Enforcement renders ZooKeeper Canary unable to connect

7.3.2

When ZooKeeper authentication enforcement was enabled, the Cloudera Manager Service Monitor's ZooKeeper Canary could not authenticate to ZooKeeper. This caused the canary checks to fail and blocked successful installation and monitoring in secured environments.

This issue is now resolved and Kerberos-based authentication is now enabled for the ZooKeeper Canary in the Service Monitor. The canary can now connect and run health checks correctly when ZooKeeper authentication is enforced.

OPSAPS-57019: No option to enforce client authentication

7.3.2

Previously, ZooKeeper did not provide an option in Cloudera Manager to enforce client authentication, limiting the ability to meet stricter security and compliance requirements.

This issue is now resolved and a ZooKeeper configuration option added in Cloudera Manager allows administrators to require client authentication.

OPSAPS-76344: ZooKeeper service canary failed with a certificate exception

7.3.2

Previously, in Cloudera Manager 7.13.2.0, ZooKeeper service canary failed due to a java.security.cert.CertificateException error. This issue is now resolved.

OPSAPS-75265, OPSAPS-75418: Mismatch in ZooKeeper component version

Previously, the ZooKeeper component version displayed and used by Cloudera Manager was not aligned with the updated ZooKeeper rebase.This issue is now fixed and Cloudera Manager now uses the correct ZooKeeper version.

Apache patch information:

- ZOOKEEPER-2590
- ZOOKEEPER-2623
- ZOOKEEPER-3100
- ZOOKEEPER-4026
- ZOOKEEPER-4236
- ZOOKEEPER-4293
- ZOOKEEPER-4293
- ZOOKEEPER-4393
- ZOOKEEPER-4604
- ZOOKEEPER-4604
- ZOOKEEPER-4728
- ZOOKEEPER-4758

- ZOOKEEPER-4762
- ZOOKEEPER-4787
- ZOOKEEPER-4787
- ZOOKEEPER-4787
- ZOOKEEPER-4843
- ZOOKEEPER-4846
- ZOOKEEPER-4876
- ZOOKEEPER-4886
- ZOOKEEPER-4889
- ZOOKEEPER-4900
- ZOOKEEPER-4909
- ZOOKEEPER-4919
- ZOOKEEPER-4933
- ZOOKEEPER-3731
- ZOOKEEPER-3860
- ZOOKEEPER-4240
- ZOOKEEPER-4415
- ZOOKEEPER-4546
- ZOOKEEPER-4753
- ZOOKEEPER-4764
- ZOOKEEPER-4799
- ZOOKEEPER-4850
- ZOOKEEPER-4860
- ZOOKEEPER-4906

Known Issues In Cloudera Runtime 7.3.2

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

Known Issues in Apache Atlas

Known issues and technical limitations for Atlas are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

CDPD-110585: Fix Janusgraph issues for jdk21 compatibility (cdh_main)

7.3.2.20000

Atlas JanusGraph integration fails on JDK 21 because reflective access to modify JanusGraph store manager internals throws an IllegalAccessException due to stricter field access restrictions.

None

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

CDPD-102546: Rename scenarios are excluded from parallel historical entity resolution

7.3.2 or higher versions

7.3.2.10000

In the distributed notification processing framework, rename scenarios during historical entity resolution are not supported. RENAME events require complex mapping tables and multi-step fallback queries that introduce significant performance overhead. Consequently, lineage adjustments

following a rename operation might exhibit discrepancies if the messages are processed out of sequence across parallel topics.

None

CDPD-94210: Lineage linkage failures in entity delete-and-recreate scenarios

7.3.2.10000 or higher versions

In the distributed notification processing framework, a critical misalignment can occur when an entity is deleted and immediately recreated (assigning a new GUID). Lineage messages for the recreated entity are often consumed before metadata for that entity is fully available or committed in the Atlas repository, which can cause lineage linkage to attach incorrectly.

None

CDPD-96101: import-hbase.sh fails with NoSuchMethodError when connecting to HBase

7.3.2.100 through 7.3.2.100 and 7.3.2.10000

7.3.2.100 and 7.3.2.10000

The Atlas HBase metadata import utility (import-hbase.sh) can fail with a NoSuchMethodError when it creates an HBase connection, for example: org.apache.hadoop.hbase.shaded.org.apache.zookeeper.client.ZKClientConfig org.apache.hadoop.hbase.zookeeper.ZKConfig.getZKClientConfig(org.apache.hadoop.conf.Configuration). The failure is caused by incompatible HBase client JARs on the import-hbase.sh classpath. Class loading can mix hbase-shaded-mapreduce or hbase-shaded-client classes with hbase-common, which produces the method signature mismatch. This issue has been observed after upgrading to JDK 17.

None

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

CDPD-109807: Async import queue permanently blocked after JanusGraph transaction failure

7.3.2.100 through 7.3.2.20000

If duplicate asynchronous metadata import requests for the same import arrive close together, they can cause a JanusGraph lock conflict, and the transaction commit fails during request registration. Atlas does not release the asynchronous import semaphore and does not update the import request status, so the request stays in the WAITING state. Atlas queues subsequent asynchronous import requests behind the blocked request and does not process them.

Restart the Atlas Server. After the restart, the import queue recovers and processes the queued import requests.

Known issues identified in Cloudera Runtime 7.3.2

CDPD-91624: Performance degradation: Atlas Replication is 3-4x slower in 7.3.2 compared to 7.1.9

7.3.2 through 7.3.2.100

7.3.2.100

A performance degradation has been observed in the Atlas table replication process after upgrading from version 7.1.9 to 7.3.2. The average time to replicate a single table has increased by approximately 300-400%. In version 7.3.2, the replication takes 3-4 seconds per table, compared to approximately 0.94 seconds per table in version 7.1.9 (calculated from 0.11s export + 0.83s import). This significant slowdown affects the time required to complete Atlas replication operations for large-scale table replication, with tests showing that only 24,000 out of 100,000 tables were replicated in more than 40 hours.

None

CDPD-92009: Unable to import Hive tables using the import script in Public Cloud

7.3.2 and its SPs and CHFs

7.3.2.10000

When you import Hive metadata using the import utility in Public Cloud, mandatory-attribute validation for shell entities can fail and block the import. The import stops with an error such as hive_table.name: mandatory attribute value missing in type Asset (ATLAS-404-00-007), even when the shell entity is expected to be populated later during import.

None

OPSAPS-75853: Composite replication (Atlas + Iceberg) history entries recorded with "Partial success"

7.3.2 and its SPs and CHFs

When running composite replication operations that include both Atlas metadata and Iceberg data replication, the replication history entries are recorded with a status of Partial success even when the replication completes successfully. This can cause confusion about the actual status of replication operations. The generic multipart upload in Atlas is facing timeout issues while uploading a file of size 291MB.

The workaround is to upload the files to the server and try the async import process.

CDPD-96774: Entities not created in Atlas when Hive table created on top of OFS bucket/volume

7.3.2 and its SPs and CHFs

7.3.2.10000

When a Hive table is created with data stored directly on an Ozone File System (OFS) bucket or volume (without a key path), the corresponding entities are not created in Apache Atlas. Only Hive tables created with OFS key-based paths (for example, ofs://serviceId/volume/bucket/key) are properly tracked in Atlas. Tables pointing directly to buckets (for example, ofs://serviceId/volume/bucket) or volumes (for example, ofs://serviceId/volume) do not generate Atlas entities, preventing metadata tracking and governance for these storage configurations.

Apache JIRA: ATLAS-5034

None

CDPD-95669: React UI: Entity type clicked in Basic Search results redirects to Home page instead of triggering type search

7.3.2 and its SPs and CHFs

7.3.2.10000

In the React-based Atlas user interface, when you perform a basic search that returns multiple entities and click an entity type in the Type column of the search results table, the interface incorrectly redirects to the Home page instead of remaining on the Basic Search page and triggering a new search filtered by that entity type. This navigation behavior breaks the expected workflow and makes it difficult to refine searches based on entity types directly from the search results.

Use the Advanced Search feature to filter results by entity type, or manually enter the type filter in the Basic Search interface before performing the search.

Apache JIRA: ATLAS-5172

CDPD-95658: React UI: isIncomplete Property not visible in React UI for shell entities

7.3.2 and its SPs and CHFs

7.3.2.10000

When viewing shell entities in the React user interface, the isIncomplete property is not displayed, even though the same entity shows this property correctly in the Classic user interface. This inconsistency prevents users from identifying incomplete entities when using the React interface. Shell entities are placeholder entities created with minimal metadata, and the isIncomplete property indicates that these entities require additional information to be complete. The React user interface now displays the isIncomplete property consistently with the Classic user interface, allowing users to identify shell entities regardless of which interface they use.

None

Apache JIRA: ATLAS-5170

CDPD-95659: Atlas UI creates multi-valued Business Metadata with SINGLE cardinality

7.3.2 and its SPs and CHFs

7.3.2.10000

When you create or edit business metadata in the Atlas user interface, the API request always sets SINGLE cardinality for attributes, regardless of whether you select Enable multiple values. Multi-valued business metadata attributes cannot be defined correctly because the UI selection is not reflected in the cardinality sent to the API.

None

Apache JIRA: ATLAS-5169

CDPD-102081: React UI redirects to Home page upon unauthorized tag or term assignment

7.3.2 and its SPs and CHFs

7.3.2.10000

In the React-based Atlas user interface, if Ranger denies permission to add a classification or glossary term, the interface redirects you to the Home page instead of displaying an authorization error message. Additionally, the Business Metadata edit window incorrectly displays a remove control (x) for existing applicable types, which cannot be removed once added. Furthermore, the Business Metadata import window truncates long file names, making it difficult to verify the selected file before uploading.

None

Apache JIRA: ATLAS-5275

CDPD-63922: RowTooBigException when reading and writing data to HBase

7.3.2 and its SPs and CHFs

When Atlas performs operations that read or write large amounts of data to HBase through JanusGraph, a RowTooBigException can occur if the row size exceeds the configured maximum row size limit (hbase.table.max.rowsize). This exception has been observed in multiple escalations during operations such as patch application, advanced search, and entity creation. The error indicates that the row data being processed is larger than the allowed maximum (typically 1GB or 2GB depending on configuration). While increasing the hbase.table.max.rowsize property can help in some cases, it may also cause out-of-memory errors on the HBase side. The root cause is related to how JanusGraph stores data in HBase and requires deeper investigation of Atlas data storage patterns.

As a temporary workaround, you can increase the hbase.table.max.rowsize property value. However, be aware that significantly increasing this value may lead to out-of-memory errors in HBase. Monitor HBase memory usage after making this change.

Apache JIRA: ATLAS-5426

CDPD-94690: React UI: Classification searches parent classification instead of child

7.3.2 and its SPs and CHFs

7.3.2.10000

When using the React user interface to search for entities by classification, clicking the Search option on a sub-classification incorrectly triggers a search for the parent classification instead of the selected sub-classification. This causes users to see search results for the wrong classification hierarchy level, making it difficult to find entities tagged with specific sub-classifications. The search now correctly uses the selected sub-classification, returning accurate results that match the user's intended search scope.

None

Apache JIRA: ATLAS-5168

CDPD-92014: React UI: Fix UI layout issue for tables when no records are present

7.3.2 and its SPs and CHFs

7.3.2.10000

When viewing tables in the React UI Search and Business Metadata sections with no records, the table headers appear visually constrained or misaligned. This layout inconsistency affects the user interface appearance and makes empty result screens look unprofessional. The table headers now display with proper alignment and layout even when no records are present, ensuring a consistent and polished user experience across all data states.

None

Apache JIRA: ATLAS-5151

CDPD-91961: Import failure due to missing support for handling multiple relationship types under a single attribute

7.3.2 and its SPs and CHFs

7.3.2.10000

When importing entities where a single relationship attribute (such as hive_db.tables) contains a mix of different entity types (for example, both hive_table and iceberg_table references), the import operation may fail or produce incomplete relationship graphs. Atlas incorrectly derives only one relationship type for the entire collection, typically hive_table_db, which causes relationship validation to fail for other entity types like iceberg_table. This data-dependent issue is difficult to diagnose and can result in import errors or missing Iceberg table relationships. Atlas logs may show errors about invalid relationship end types or invalid relationship definitions for specific elements within a relationship attribute.

None

Apache JIRA: ATLAS-5156

CDPD-87879: Moving an Iceberg table from one database to another database is creating incomplete entity

7.3.2 and its SPs and CHFs

When renaming an Iceberg table to move it from one database to another using the ALTER TABLE statement with RENAME TO clause (for example, ALTER TABLE test_share.iceberg_table RENAME TO default.iceberg_table), an incomplete entity is created in the target database. This occurs because the Hive Metastore hook incorrectly classifies the operation as an ADDPROPS (add properties) operation instead of a RENAME operation for Iceberg tables. This misclassification causes the HMS hook to process the event incorrectly, passing only the DO_NOT_UPDATE_STATS = true property while omitting critical properties such as old_table_name, alterTableOpType = RENAME, and old_db_name that are required for proper entity handling. As a result, Atlas creates an incomplete entity in the target database rather than properly renaming the existing entity.

None

CDPD-99217: Incomplete Metadata Display for Ozone Buckets in Atlas Search Results

7.3.2 and its SPs and CHFs

While searching for certain Ozone buckets in the Atlas UI, the results page loads, but key details such as Properties, Classifications, Relationships, and Audits are not displayed. The issue stems from the large size of the response payload, specifically the relationships section, which exceeds 500MB in the JSON response. This overwhelms the UI, preventing it from rendering associated metadata.

None

7.3.2.10000

Apache JIRA: ATLAS-5184

CDPD-91959: Entities not created in Atlas when hive table created on top of a OFS bucket/volume

7.3.2 and its SPs and CHFs

In Atlas, key based tables are created properly. However, bucket/volume based tables do not show up in Atlas. When trying to create hive tables on top of ozone OFS buckets/volumes, the entities are not created.

None

Apache JIRA: ATLAS-5034

CDPD-89476: need to enhance the repair_index.py script

7.3.2 and its SPs and CHFs

When executing the repair_index.py script provided by Cloudera to reindex Atlas metadata, the script fails to run under Java 17, which is now commonly used in customer environments. Customers upgrading their OS or JVM to OpenJDK 17 encounter errors related to java.lang.reflect.InaccessibleObjectException. Even after bypassing this with additional JVM flags, the script later fails with Keystore was tampered with, or password was incorrect. This indicates potential mismatches with the Java version, keytool compatibility, or the way the script selects runtime process directories and handles credentials.

Short-term workaround:

- Add the following JVM options in repair_index.py to bypass JPMS restrictions in Java 17:

```
--add-opens=java.base/java.lang=ALL-UNNAMED \
                --add-opens=java.base/java.lang.reflect=ALL-UNNAMED
```

- Manually verify and specify the correct conf and keystore directory to avoid loading stale credentials.
- Run the script under Java 8 or Java 11 if possible.

CDPD-88306: Atlas UI: Iceberg tables are not visible in the Relationships tab

7.3.2 through 7.3.2.10000

7.3.2.10000

On the Atlas user interface, the **Relationship** tab for a database entity is not displaying Iceberg tables when both Hive and Iceberg tables are present in the same database. While Hive tables are visible as expected, Iceberg tables are missing from the visualization. If only Iceberg tables exist in the database, they are displayed correctly in the **Relationships** tab.

None

CDPD-78683: Standalone Atlas replication fails when unexpected inputs are provided

7.3.2 and its SPs and CHFs

When a standalone Atlas replication policy is created for NULL or hive_db entity types, and no Hive or Iceberg tables are created, the replication fails with a generic exception (Atlas Export operation has failed. Please check role logs for more info.) instead of either skipping the operation (NOOP) or providing a descriptive error message about the problematic inputs. This generic error message can be misleading.

None

CDPD-76686: Concurrent Ingestion: doesn't commit the kafka offset

7.3.2 and its SPs and CHFs

When concurrent processing is enabled (atlas.notifications.concurrent=true), Atlas processes messages efficiently but may fail to commit the Kafka offset. This results in the current offset getting stuck for an extended period, despite the messages being processed successfully.

None

CDPD-77640: Inconsistent results from global search due to case sensitivity

7.3.2 and its SPs and CHFs

Global search in Atlas produces inconsistent results based on the case of the search term. Searching with a lowercase term returns the expected entities, whereas searching with the same term beginning with an uppercase character does not return the same results. This issue is related to case-sensitivity handling in the search layer, likely in indexing or query processing.

None

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-80868: Basic search API returns wrong response for attribute-based search with long values

7.1.9 through 7.1.9 SP2 CHF1, 7.3.1, 7.3.1.100, 7.3.1.200, 7.3.1.300, 7.3.1.400, 7.3.1.500, 7.3.1.600, 7.3.1.706, 7.3.1.900

Basic search using attribute filters, such as qualifiedName, returns no results when the filter values exceed the default maximum token length of 255 characters configured in Solr. This issue affects the startsWith, endsWith, and contains operators.

Use the Advanced Search (DSL) or direct entity lookup by GUID, which are not affected by this Solr token length limitation.

Apache Jira: ATLAS-5032

CDPD-103032: Business metadata with an array attribute type is slow to delete

, 7.1.9 and its SPs and CHFs, 7.3.1 through 7.3.1.1000, 7.3.2 through 7.3.2.10000

7.3.2.10000

When you delete a business metadata typedef that includes an attribute with the array data type, Apache Atlas can take a long time because it scans Solr for entity references even though array attribute values are not Solr-indexed.

None.

Apache JIRAATLAS-4988

CDPD-87937: Incomplete metadata display for Ozone buckets in Atlas search results

7.1.9 CHF4 through 7.1.9 CHF17 7.3.1.500, 7.3.1.600, 7.3.1.706, 7.3.1.900, 7.3.2 and its SPs and CHFs

7.1.9.20000, 7.3.2.100, 7.3.2.10000

For some Apache Ozone buckets that you open from search results in Apache Atlas, the entity detail page does not display Properties, Classifications, Relationships, and Audits. This happens because the relationship payload is too large and the user interface becomes unresponsive.

None.

Apache Jira: ATLAS-5184

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

7.3.2 and its SPs and CHFs, 7.3.1.700 through 7.3.1.100, 7.1.9 SP1 CHF15 through 7.1.9 SP1 CHF4

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-69150: Unable to add labels or user defined properties in Japanese

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Adding Japanese labels or user defined properties results in the error message: “Invalid label: ###, label should contain alphanumeric characters, _ or -”

None

CDPD-69279: Quick search does not return entities when using Japanese or Chinese characters to search properties

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Entities, such as hdfs_path, are not returned by the search when their indexable properties are searched using partial search terms made of Japanese or Chinese characters. Only exact matches return results when searching the indexed properties made of Japanese or Chinese characters.

None

CDPD-56085: [Impala Iceberg] LOAD DATA INPATH to Iceberg_table creates a temporary hive_table with name <iceberg_table_name>_tmp* and then marks it as DELETED in Atlas

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Running a query like LOAD DATA INPATH to iceberg_table, creates a temporary hive table with name <iceberg_table_name>_tmp* and then marks it as DELETED in Atlas. So in Atlas, a deleted entity is created corresponding to the temporary table <iceberg_table_name>_tmp*.

Tag added to the File system (HDFS) entity will not be propagated to the Iceberg table, user has to manually add to the Iceberg table, since the tag propagation is broken due to the deleted table in the flow.

CDPD-55301: The ddlQueries and ALTERTABLE_* lineage are missing for Spark tables created using spark3-shell

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

The ddlQueries and outputFromProcesses (lineage) is missing for the alter queries.

Apache JIRA: ATLAS-4747

CDPD-75994: Post DL regular upgrade (non ZDU) to 7.3.1, "Exception in getKafkaConsumer ,WakeupException: null" is seen

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

After the data lake is upgraded to 7.3.1, sometimes Atlas Hook does not function when Apache Atlas and Apache Kafka are started at the same time, thus Atlas is unable to connect to Kafka while Kafka is still being set up. Atlas performs only three attempts.

Restart the cluster, after the upgrade to trigger to reconnect to Apache Kafka. The Kafka consumer creation should be retried if the Kafka service is unavailable during Atlas startup.

CDPD-76035: Resource lookup for Atlas service is failing

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Once the Atlas configuration snippet atlas.authentication.method.file is enabled and a classification is created, these do not synchronize correctly to the Type Category resource field setting of Apache Ranger. The newly created classification won't be able to be selected as the Type Name.

CDPD-69140: Incremental export: When an entity has tag propagated and is exported, the tag is not propagated to it in the export

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When creating tables with multiple levels of depth, a tag applied to the first table will be propagated along the lineage. After exporting the last table, it will not have the propagated tag.

Apache JIRA: ATLAS-4889

CDPD-77128: Log out message for Atlas needs to be more informative on Knox enabled cluster

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

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

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

CDPD-83398: REMOVE_LEGACY_REF_ATTRIBUTES Patch Skipped During Startup, Causing Errors While Processing DML Events

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHF

The REMOVE_LEGACY_REF_ATTRIBUTES patch is skipped during Atlas patch processing, which leads to an error when Data Manipulation Language (DML) events (e.g., "insert data") are processed, events that are typically ignored. This is observed in environments where DML event filtering is not in place or the patch is not applied correctly. As a result, Atlas attempts to process these events and encounters errors related to missing or outdated attribute handling.

CDPD-88230: Bulk import operations currently do not construct lineage as expected

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Bulk import operations do not create lineage consistently, even when the entities are imported successfully.

CDPD-88302: Bulk import operations currently do not create audits for the imported entities

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Bulk import operations do not create audits for the imported entities.

CDPD-7982: HBase bridge stops at HBase table with deleted column family

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Bridge importing metadata from HBase fails when it encounters an HBase table for which a column family was previously dropped. The error indicates:

```
Metadata service API
            org.apache.atlas.AtlasClientV2$API_V2@58112bc4 failed
 with status 404 (Not Found)
            Response Body
              ({""errorCode"":""ATLAS-404-00-007"",""errorMessa
ge"":""Invalid instance
            creation/updation parameters passed :
              hbase_column_family.table: mandatory attribute v
alue missing in type
            hbase_column_family""})
```

None

Apache JIRA: ATLAS-3551

CDPD-6675: Irregular qualifiedName format for Azure storage

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

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-922: IsUnique relationship attribute not honored

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

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

None

CDPD-1664: Guest users are redirected incorrectly

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Authenticated users logging in to Atlas are redirected to the Cloudera 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-1823: Queries with ? wildcard return unexpected results

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

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 Cloudera.

None

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

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

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

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

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

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

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-3160: Hive lineage missing for INSERT OVERWRITE queries

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

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-3208: Table alias values are not found in search

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

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

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

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-4762: Spark metadata order may affect lineage

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

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-10574: Suggestion order doesn't match search weights

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

At this time, the order of search suggestions does not honor the search weight for attributes.

None

Apache JIRA: ATLAS-3081

CDPD-10576: Deleted Business Metadata attributes appear in Search Suggestions

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Atlas search suggestions continue to show Business Metadata attributes even if the attributes have been deleted.

None

Apache JIRA: ATLAS-3739

CDPD-11692: Navigator table creation time not converted to Atlas

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

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

None

Apache JIRA: ATLAS-3739

CDPD-11940: Database audit record misses table delete

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

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

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

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-55122: Any user with ssh access can view the downloaded results

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

When a user triggers a download from the Atlas user interface, the downloaded data can be viewed by any other user who has SSH access to the host. This occurs because the downloaded files are not sufficiently restricted by file permissions on the local filesystem, allowing unintended access to potentially sensitive data.

None

CDPD-67112: Import transforms do not work as expected when replacing a string which already has ":"

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

The character ":" is not supported in path replacements. The import succeeds but location remains unchanged. The character " :" must be avoided.

None

Apache JIRA: ATLAS-4834

CDPD-56590: Create table "like" from Iceberg table creates a hive_table instead of iceberg_table

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

By default, for tables created using the "like" command, lineage is not generated in Atlas. The destination like table should be of the same type as source table. Instead an iceberg_table for source and hive_table for destination are getting created.

CDPD-68191: Suggestions do not return the correct results when searching multiple Chinese characters

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

Free text search does not return results when searching Chinese phrases made of multiple characters. Partial searches return the correct results.

None.

CDPD-77434: Search suggestions may fail to return results when using complete qualified names in RAZ-enabled environments

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When performing searches in the Atlas UI or through the REST API using a complete qualified name as the prefix string (for example, searching for an Iceberg column using its full qualified name), the suggestions API sometimes fails to return the expected entity in the suggestions list, even though the entity exists and would be returned in the actual search results. This issue occurs consistently in RAZ-enabled AWS clusters but has not been observed in IDB (internal data bus) or Cloudera on prem clusters. The behavior is intermittent: the REST API /api/atlas/v2/search/suggestions?prefixString=<complete_qualifiedName> may return an empty suggestions array ({"suggestions": []}), while the same entity appears correctly when searched through the user interface. This inconsistency affects user experience and makes entity discovery less reliable in RAZ-enabled environments.

None

CDPD-81499: Basic search performance degradation when filtering on boolean attributes

7.3.2 and its SPs and CHFs, 7.3.100 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When performing basic searches with filters on boolean attributes (for example, isDisposed = true), search operations may experience significant performance degradation, particularly with large datasets. This occurs because boolean datatype attributes are not indexed in Solr and require in-memory filtering in Atlas. When the search API attempts to return the requested number of results (for example, limit of 25), it repeatedly queries Solr for active entities and then filters them in Atlas based on boolean attribute values. This causes Atlas to traverse through potentially large numbers of entities until it finds enough matches, resulting in increased search times as data volume and query offset increase. Atlas logs may display warnings such as: not using index-search for attribute EntityType.booleanAttribute; might cause poor performance.

None

CDPD-69317: Export/Import API: When tag attributes types get modified between the exports, causes issues in import

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When using the Export/Import API, if a tag's attribute type is modified (for example, from string to double) between incremental exports, the import operation on the target cluster fails due to the difference in attribute types.

None

CDPD-65906: Export: When relationship is updated to block tag , the imported lineage doesn't have the updated relationship

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its SPs and CHFs

When using the Export/Import API to migrate data between clusters, if a relationship between entities is modified to block tag propagation (for example, setting propagation from table to process as 'None') and a tag is subsequently added, the tag correctly does not propagate on the source cluster. However, after exporting to the target cluster, the relationship on the target cluster does not reflect the blocked propagation, and the tag incorrectly propagates to the related entities.

None

Apache JIRA: ATLAS-4833

Known Issues in Apache Avro

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

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known Issues in Cloud Connectors

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

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified before Cloudera Runtime 7.3.2

There are no known issues carried forward that affect Cloud Connectors in Cloudera Runtime 7.3.2.

Known Issues in Cruise Control

Known issues and technical limitations for Cruise Control are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-44676: Rebalancing with Cruise Control does not work if the metric reporter fails to report the CPU usage metric

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

If the CPU usage metric is not reported, the numValidWindows in Cruise Control will be 0 and proposal generation as well as partition rebalancing will not work. If this issue is present, the following message will be included in the Kafka logs:

```
WARN com.linkedin.kafka.cruisecontrol.metricsreporter.CruiseCont
rolMetricsReporter:
       [CruiseControlMetricsReporterRunner]: Failed reporting CPU
 util.
```

```
java.io.IOException: Java Virtual Machine recent CPU usage is not
 available.
```

This issue is only known to affect Kafka broker hosts that have the following specifications:

- CPU: Intel(R) Xeon(R) CPU E5-2699 v4 @ 2.20GHz
- OS: Linux 4.18.5-1.el7.elrepo.x86_64 #1 SMP Fri Aug 24 11:35:05 EDT 2018 x86_64
- Java version: 8-18

Move the broker to a different machine where the CPU is different. This can be done by moving the host to a different cluster. For more information, see Moving a Host Between Clusters

Note: Cluster nodes affected by this issue are not displayed as unhealthy.

Known Issues in Apache HBase

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

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

Client-side TLS hostname validation issue in Apache Thrift (libthrift)

7.3.2.20000

A TLS hostname validation issue exists in libthrift versions lower than 0.23.0. The issue does not allow remote compromise of your CDP HBase server, but external client applications that connect to the HBase Thrift gateway over TLS might be vulnerable to Man-in-the-Middle (MitM) traffic interception if an attacker redirects network traffic.

Update libthrift to 0.23.0 or higher versions in your client applications. Version 0.23.0 is fully wire-compatible with the CDP HBase Thrift server. If an immediate client library update is not possible, perform one of the following actions:

- Enforce Kerberos (SASL) authentication on the HBase Thrift gateway by using Cloudera Manager.
- Configure client socket settings to explicitly perform TLS hostname verification.

For more information, see CVE-2026-43869.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

HBase replication between TLS-enabled Cloudera Operational Database clusters fails due to hostname verification

7.3.2.10000

When you set up reverse HBase replication between two TLS-enabled Cloudera Operational Database clusters, replication fails if reverse DNS resolution is not configured between the clusters. During mutual TLS (mTLS) authentication, the source cluster sends its certificate, which contains the hostname. The target cluster validates the certificate against its truststore and then verifies that the client connection IP address matches the hostname in the certificate. Because reverse DNS resolution does not work between separate Cloudera Operational Database clusters by default, this verification check fails and the connection is rejected.

This issue occurs when all of the following conditions are true:

- Both clusters have TLS enabled.
- Hostname verification is enabled (hbase.client.netty.tls.verify.server.hostname and hbase.server.netty.tls.verify.client.hostname are set to true).
- hbase.server.netty.tls.client.auth.mode is set to WANT or NEED.
- TLS certificates are issued without an IP Subject Alternative Name (SAN).
- Reverse DNS resolution is not configured between the clusters.
- Mutual TLS trust is established between the clusters.

These are the default settings for a fresh installation of Cloudera Runtime 7.3.2.10000.

To bypass the client hostname verification check, add the hbase.server.netty.tls.verify.client.hostname = false property to the HBase Service Advanced Configuration Snippet (Safety Valve) for hbase-site.xml using Cloudera Manager.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified before Cloudera Runtime 7.3.2

There are no known issues carried forward that affect HBase in Cloudera Runtime 7.3.2.

Known Issues in HDFS

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

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known Issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known Issues Identified Before Cloudera Runtime 7.3.2

CDPD-101357: Using Windows NFS Client to connect to HDFS NFS gateway is not supported

7.1.7 and its SPs and CHFs, 7.1.8 and its SPs and CHFs, 7.1.9 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, and 7.3.2 and its SPs and CHFs

When attempting to copy files larger than 1.5 MB from a Windows machine to an NFS mount point, the operation fails with the Can't read from the source file or disk error. However, files smaller than 1 MB are copied successfully.

Do the following steps:

1. **Windows to Linux Folder Sharing**: Share a folder from Windows and mount it on the same Linux host, then perform the file copy operation within the Linux environment.
2. **Linux File System Sharing**: Share the Linux file system and mount it on the Windows machine, then perform the file copy operation within the Linux environment.

Known Issues in Apache Hive

Known issues and technical limitations for Hive are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

CDPD-111552: Altering an Iceberg table column to an unsupported type returns an unclear internal error

7.3.2.20000

When you used an ALTER TABLE ... CHANGE COLUMN statement to change the data type of a column in a Hive table stored by Iceberg to a type that is not a valid promotion, for example changing a BOOLEAN column to STRING, or narrowing a BIGINT column to INT, Hive did not return a clear, user-friendly error message. Instead, Hive returned an execution error that exposed a raw, unhandled internal stack trace, for example Execution Error, return code 40013 from org.apache.hadoop.hive.ql.ddl.DDLTask. Unable to alter table. java.lang.IllegalArgumentException: Cannot change column type. The same behavior occurred when the statement was run from Beeline. Valid Iceberg type promotions, such as changing an INT column to BIGINT, continued to work as expected.

None.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2

CDPD-98415: Hive ACID incremental replication fails with "MetaException: invalid write id"

7.3.2.0

During Hive ACID incremental replication, the replication load task fails with a metastore exception similar to the following:

```
Caused by: org.apache.hadoop.hive.metastore.api.MetaException: i
nvalid write id <ID> for table <db_name>.<table_name>
```

This issue occurs when processing a large volume of incremental replication events (such as bulk DDLs, partition operations, or TRUNCATE TABLE queries).

When the load builder batches events across multiple task iterations (controlled by hive.repl.approx.max.load.tasks), proactive transaction cleanup logic (CLEAR_DANGLING_TXNS) executes prematurely at the end of a batch iteration instead of at the end of the full incremental load operation. If an open transaction is replayed in one batch, but its corresponding commit event is scheduled in the subsequent batch, the cleanup task prematurely aborts the in-flight target transaction and invalidates its write ID.

If an upgrade to Cloudera Runtime 7.1.9 SP2 is not immediately possible, disable the proactive dangling transaction cleanup logic and enforce strict database cleanup ordering:

1. **Disable dangling transaction cleanup:**

Set hive.repl.clear.dangling.txns.on.target to false globally in Cloudera Manager under Hive Service Advanced Configuration Snippet (Safety Valve) or explicitly in both REPL DUMP and REPL LOAD commands:

a. Source Cluster (REPL DUMP):

```
REPL DUMP <db_name> WITH (
  'hive.repl.replica.external.table.base.dir'='hdfs://<target-
hdfs-host>:8020/',
  'hive.repl.rootdir'='hdfs://<target-hdfs-host>:8020/wareho
use/tablespace/staging/<db_name>',
  'hive.repl.run.data.copy.tasks.on.target'='true',
  'hive.repl.clear.dangling.txns.on.target'='false'
);
```

b. Target Cluster (REPL LOAD):

```
REPL LOAD <db_name> INTO <db_name> WITH (
  'hive.repl.replica.external.table.base.dir'='hdfs://<target-
hdfs-host>:8020/',
  'hive.repl.rootdir'='hdfs://<target-hdfs-host>:8020/warehou
se/tablespace/staging/<db_name>',
  'hive.repl.run.data.copy.tasks.on.target'='true',
  'hive.repl.clear.dangling.txns.on.target'='false'
);
```

2. **Enforce database cleanup order during re-bootstrap:**

When hive.repl.clear.dangling.txns.on.target is set to false, administrators must follow a strict order during re-bootstrap operations to prevent orphan entries in the Hive Metastore:

a. Delete the Replication Policy First: Always remove the replication schedule or policy before deleting the database.

b. Drop the Target Database Second: Dropping the database before the policy removes database metadata while leaving stale entries in REPL_TXN_MAP.

c. Verify Metastore State: Before initiating a re-bootstrap load, run the following SQL query directly on the Target Hive Metastore database to verify no lingering entries exist for the policy:

```
SELECT * FROM REPL_TXN_MAP WHERE RTM_REPL_POLICY = '<db_name
>.*';
```

Apache Jira: HIVE-29459

CDPD-73781: High operational costs from automatic partition discovery on cloud storage

7.3.1 and its SPs and CHFs, 7.2.18 and its SPs

In cloud environments, the Partition Management Task (PMT) could significantly increase operational costs and performance overhead for large tables with many partitions.

To address this issue, you can perform the following:

- Increase the value of the metastore.partition.management.task.frequency property to 86400 seconds.
- Disable partition discovery manually for specific tables incurring high costs by using the discover.partitions table property.

CDPD-93090: Failure when inserting data into partitions located on a different filesystem than the table

7.3.2

Inserting data into a table partition results in a MoveTask execution error if the partition is stored in a different filesystem than the parent table. For example, if you create an external table in S3 but add a partition with an HDFS location, the `INSERT` operation fails with a Wrong FS error. Although the data is physically inserted into the correct location, the query returns a non-zero exit code and displays an error message.

None

CDPD-92560: Partition insertion errors across multiple buckets in RAZ enabled clusters

7.3.2

Inserting data into a table partition results in a MoveTask execution error when the partition is located in a different S3 bucket than the base table. In Ranger Access Control (RAZ) enabled Private Cloud clusters, the `INSERT` operation fails with a Wrong FS error, even though the data is physically written to the specified partition location.

None

CDPD-74192: Hive SSL certificate error during large record copies

7.3.1 and its SPs and CHFs, 7.2.18.300, 7.2.18.400

When you copy a large number of records (approximately 1 million or more) using Hive, the operation might fail with a

```
javax.net.ssl.SSLException: org.bouncycastle.tls.TlsFatalAlert:
 certificate_unknown(46)
```

error. This issue occurs during the job commit phase when Hive attempts to execute HTTP requests to S3 for file operations such as copyFile.

None

CDPD-68096: Unable to set S3 credentials at the session level in Hive

7.3.2

Accessing S3 buckets from Hive on on premises results in a MetaException or AccessDeniedException when you attempt to create external tables by using session-level credentials. While Spark and other Hadoop jobs allow you to define fs.s3a.access.key and fs.s3a.secret.key at the session level, Hive Metastore (HMS) does not honor these session-specific parameters during path validation.

To access S3 buckets, you must configure the S3 credentials at the global level in the Hive or Hadoop cluster configuration (such as core-site.xml) or use a supported credential provider like the Hadoop Credential Shell to store keys securely. Setting these parameters at the session level is currently not supported for Hive table creation.

HIVE-16913

Known issues identified before Cloudera Runtime 7.3.2

DWX-22436: DL upgrade recovery fails due to Metastore schema incompatibility

7.3.1.600

When attempting a Data Lake (DL) upgrade recovery from version 7.2.18.1100 to Cloudera Runtime 7.3.1.500, the process fails because the Hive Metastore schema versions are incompatible. The error indicates a mismatch between the Hive version (3.1.3000.7.3.1.500-182) and the database schema version (3.1.3000.7.2.18.0-Update2). This blocks Data Lake recovery if an upgrade fails, impacting customers.

Before you initiate the recovery process, manually update the Hive Metastore schema to match the target version by using the `schematool` utility.

1. Obtain the Hive database password: Run the following command to retrieve the password from the pillar configuration:

```
cat /srv/pillar/postgresql/postgre.sls
```

2. Back up the existing configuration: Move the current configuration directory to a backup location:

```
mv /etc/hive/conf /etc/hive/conf_backup
                  mkdir /etc/hive/conf
```

3. Prepare the temporary configuration: Copy the process files to the new configuration directory:

```
scp /var/run/cloudera-scm-agent/process/<process-id>-hive-me
tastore-create-tables/* /etc/hive/conf/
```

4. Update the connection password: Open the /etc/hive/conf/hive-site.xml file and perform the following modifications:
   - Set the javax.jdo.option.ConnectionPassword property to your Hive database password.
   - Comment out the hadoop.security.credential.provider.path property.
5. Run the schema upgrade tool: Execute the `schematool` to synchronize the version:

```
/opt/cloudera/parcels/CDH/lib/hive/bin/schematool -dbType po
stgres -initOrUpgradeSchema --verbose
```

6. Restore the original configuration: Remove the temporary directory and restore your backup:

```
rm -rf /etc/hive/conf
                  mv /etc/hive/conf_backup /etc/hive/conf
```

7. Restart the cluster: Restart the services to initialize the Hive Metastore with the updated schema.

CDPD-77738: Atlas hook authorization issue causing HiveCreateSysDb timeout

7.1.9 SP1 CHF4, 7.1.7 SP3 CHF7, 7.3.1.100, and its higher versions

Atlas hook authorization error causes HiveCreateSysDb command to time out due to repeated retries.

None

CDPD-74680: DAG not retried after failure

7.3.1 and its higher versions

When executing a Hive query, if the ApplicationMaster container fails, Hive does not retry the DAG if the failure message contains some diagnostic information including a line break, leading to query failure (instead of retry).

None

Known Issues in Cloudera Data Explorer (Hue)

Known issues and technical limitations for Cloudera Data Explorer (Hue) are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

CDPD-119579: Data Explorer logs page redirection

7.3.2.10000, 7.3.2.20000

When attempting to access the Data Explorer logs page, the application redirects you to the SQL Editor instead of displaying the log details.

To inspect the logs, administrators must SSH into the cluster directly and review the log files from the host environment.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2

CDPD-98754: Data Explorer HDFS access issue in Data Discovery Template

7.3.2, 7.3.2.100, 7.3.2.10000, 7.3.2.20000

In the Data Discovery Template, Data Explorer cannot access HDFS. This issue occurs because Cloudera Manager assigns a standby HDFS NameNode to the webhdfs_url property instead of the active NameNode, which prevents the file browser from displaying HDFS data.

File and folder name character corruption in HttpFS — When you rename files or folders, combine files, or set custom metadata attributes by using HttpFS on port 14000, specific special characters in file and folder names are automatically converted into HTML entity encoding. For example, renaming a file to myfile&test saves the file on HDFS as myfile&test.

The following character conversions occur in HttpFS:

- & — Converts to &
- < — Converts to <
- > — Converts to >
- '— Converts to '
- " — Converts to "

For inaccessible HDFS in Data Discovery Template:

1. Log in to Cloudera Manager as an Administrator.
2. Go to Clusters HDFS to identify which NameNode is currently active.
3. Go to Clusters Data Explorer (Hue) Configuration .
4. Change the value of the webhdfs_url property from the standby NameNode to the active NameNode.
5. Click Save Changes.
6. Restart the Data Explorer service.

For HttpFS character corruption:

- Do not include &, <, >, ', or " characters in file or folder names.
- Perform file management operations by using native WebHDFS on NameNode port 20102 instead of HttpFS port 14000.

CDPD-95911: File operations fail for files with special characters in GS RAZ environments

7.3.2, 7.3.2.100, 7.3.2.10000, 7.3.2.20000

In GS RAZ environments, you cannot rename, move, or copy files that contain special characters such as the percent sign (%). When you attempt these operations, the Failed to receive a valid re sponse from RAZ error message is displayed. Additionally, operations on files containing the hash character (#) do not result in an error, but the action is not performed.

None.

CDPD-94944: File and directory names contain special characters after a rename operation

7.3.2, 7.3.2.100, 7.3.2.10000, 7.3.2.20000

In Data Explorer, file and directory names that contain special characters, such as the ampersand (&) and the apostrophe ('), are incorrectly transformed into HTML entities after a rename operation. For example, the ampersand is replaced by & and the apostrophe is replaced by '. This issue occurs during rename operations for files with names containing characters such as ~@$&()*! +'=;.txt.

None.

CDPD-95666: Endless progress bar for S3 folders without read permissions on Azure

7.3.2, 7.3.2.100

7.3.2.10000

In Azure environments, a persistent loading bar is displayed when you select a folder or file for which you do not have read permissions. This issue occurs across the file system, including S3 and HDFS. Currently, the user interface does not display an alert message to indicate that permissions are missing. Instead, the UI remains in a loading state.

None.

CDPD-93181: Incorrect Data Explorer server status reporting

7.3.2, 7.3.2.100, 7.3.2.10000, 7.3.2.20000

Data Explorer server status is incorrectly reported as Green (Healthy) in Cloudera Manager even if the process fails to start. This occurs when the mysqlclient dependency is missing, causing database migrations to fail. The Data Explorer server appears healthy because the loglistener.py sidecar process starts first and maintains a Process ID (PID). Cloudera Manager monitors this PID and reports a successful start, even though the main Data Explorer process has terminated due to the missing dependency.

1. Stop the Data Explorer service in Cloudera Manager.
2. Log in to the affected Hue host and manually terminate the loglistener.py sidecar process.
3. Remove any stale PID files located in the /tmp/ directory, for example, /tmp/hue_*.pid.
4. Ensure that the mysqlclient dependency is installed on the host before attempting to restart the service.

Known issues identified Before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-95086: Data Explorer UI fails to load due to a connection timeout error

7.3.2, 7.3.1.706

In a Cloudera Data Hub high availability (HA) environment, the Data Explorer UI can fail to load after the master0 HDFS node is deleted. Data Explorer continues to attempt connections to the deleted node because Data Explorer does not support HDFS HA. As a result, the Data Explorer file browser points to an invalid HDFS endpoint, which causes connection timeouts and renders the UI inaccessible.

You can restore access to the Data Explorer UI by updating the configured HttpFS node to an active HDFS node in Cloudera Manager and then restarting the Data Explorer service.

CDPD-92946: Data Explorer debug logging cannot be disabled

7.3.2, 7.3.1.600, 7.3.1.706

Data Explorer debug logging is always enabled and cannot be turned off.

None.

CDPD-88964: Data Explorer logs missing in Data Explorer UI

7.3.2, 7.3.1.500 through 7.3.1.706

The Data Explorer UI might display only a few lines of logs instead of the complete Data Explorer logs. This can occur due to leftover Gunicorn processes that interfere with the proper logging and display of logs within the Data Explorer interface.

1. Stop the Data Explorer service.
2. Terminate any remaining Gunicorn processes to clear hung or orphan processes that might be causing the issue. Run the following command (use sudo if not running as root):

```
# pids=$(ps -efwww | grep rungunicornserver | grep -v "grep" |
 grep "rungunicornserver"| awk '{ print $2 }') && for i in $pi
ds; do kill -9 $i ; done
```

3. Restart the Data Explorer service.

CDPD-90510: Defunct Data Explorer Gunicorn worker processes accumulate

7.3.2, 7.3.1.500 through 7.3.1.706

Data Explorer on Ubuntu 22 using Oracle Database can accumulate defunct rungunicornserver worker processes due to incomplete process termination and stale database connections. This can lead to a cluttered process table, which does not critically impact service functionality.

Periodically clean defunct worker processes by running the following command:

```
pgrep -f 'hue rungunicornserver' | xargs -r kill -9
```

OPSAPS-75134: LDAP and Kerberos dual authentication fails with HiveOnTez in HTTP Transport Mode

7.3.2, 7.3.1 and its SPs and CHFs, Cloudera Manager 7.13.1 and its SPs and CHFs

Enabling LDAP for the HiveOnTez service in a Kerberos environment with transport mode set to HTTP, Data Explorer fails to load database information . This failure occurs because Data Explorer does not support the combined KERBEROS and LDAP authentication value required by HiveOnTez. The resulting conflict prevents a successful connection to HiveServer2

- Log in to Cloudera Manager Hive-On-Tez Configuration Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml and set the following value:

```
hive.server2.authentication=LDAP,KERBEROS
```

- Go to Hue Configuration Hue Server Advanced Configuration Snippet (Safety Valve) for hive-site.xml and set the following value:

```
hive.server2.authentication=KERBEROS
```

CDPD-58978: Batch query execution using Data Explorer fails with Kerberos error

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

When you run Impala queries in a batch mode, you encounter failures with a Kerberos error even if the keytab is configured correctly. This is because submitting Impala, Sqoop, Pig, or pyspark queries in a batch mode launches a shell script Oozie job from Data Explorer and this is not supported on a secure cluster.

None. Submit the queries individually.

CDPD-59677: Unable to view Phoenix tables on the left assist in Data Explorer

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

On clusters secured with Knox, you might not be able to see Phoenix tables on the left assist that are present under the default database (that is, an empty('') database).

None.

Known Issues in Apache Iceberg

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

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

CDPD-111552: Unsupported Iceberg column type alterations in Hive return unhandled IllegalArgumentException

Using Hive to change the data type of a column on an Iceberg table to an unsupported conversion, such as BOOLEAN to STRING or a narrowing change such as BIGINT to INT results in the statement failure with an unhandled IllegalArgumentException and a Java stack trace. Hive does not return a concise, user-facing error message (for example, that the conversion is not supported). Supported Iceberg type promotions, such as INT to BIGINT, complete successfully. The same behavior occurs in Beeline.

Use only Iceberg-supported column type promotions when you alter Iceberg table schemas. Unsupported conversions are rejected. If a statement fails, use the Cannot change column type message in the error output to identify the unsupported conversion.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

CDPD-101742: Hive Iceberg snapshot expiration fails for ORC LZ4/LZO tables

Hive fails to run queries on Iceberg tables stored in ORC format with LZ4 or LZO compression when LLAP cache is enabled. The failure results in an OWN_TASK_FAILURE status and vertex failures due to an UnsupportedOperationException during ORC decompression. The snapshot expiration functionality itself works as expected and is not impacted.

For Iceberg tables, disable vectorization by setting hive.vectorized.execution.enabled=false.

CDPD-89390/CDPD-83022: Incorrect row count displayed in table metadata after compaction

7.3.2.10000, 7.3.2, 7.3.1.600, 7.3.1.500, 7.3.1.400

After running data compaction operations on large tables, the row count displayed by the DESC RIBE FORMATTED command may be inaccurate. Initially, the count may appear higher than the actual number of rows. Subsequently, after running the ANALYZE command to update table statistics, the count might then appear lower than the actual number of rows.

This issue has been observed in large tables containing a significant number of historical snapshots (exceeding 10000). All these snapshots are primarily generated through UPDATE operations.

It is important to note that this is just a metadata display issue, and there is no loss of data. The underlying table data remains complete and correct.

To obtain an accurate row count, use the SELECT COUNT(*) query.

Cloudera Lakehouse Optimizer

CDPD-119573: Evaluation API returns an error

The evaluation API returns a Duplicate evaluation request error when stale task metadata exists in the submitted state.

Delete the stale task metadata that is in the submitted state for the affected table from storage (HDFS or S3).

Known Issues identified in Cloudera Runtime 7.3.2.100 CHF1

There are no new known issues in this release.

Known Issues identified in Cloudera Runtime 7.3.2

DWX-18843: Unable to read Iceberg table from Hive Virtual Warehouse

7.3.2

If you have used Apache Flink to insert data into an Iceberg table that is created from Hive, you cannot read the Iceberg table from the Hive Virtual Warehouse.

Add the engine.hive.enabled table property through the Hive beeline and set the value to "true". You can add this table property either while creating the Iceberg table or use the ALTER TABLE statement to add the table property.

DWX-18489: Hive compaction of Iceberg tables results in a failure

7.3.2

When Cloudera Data Warehouse and Cloudera Data Hub are deployed in the same environment and use the same Hive Metastore (HMS) instance, the Cloudera Data Hub compaction workers can inadvertently pick up Iceberg compaction tasks. Since Iceberg compaction is not yet supported in the latest Cloudera Data Hub version, the compaction tasks will fail when they are processed by the Cloudera Data Hub compaction workers.

In such a scenario where both Cloudera Data Warehouse and Cloudera Data Hub share the same HMS instance and there is a requirement to run both Hive ACID and Iceberg compaction jobs, it is recommended that you use the Cloudera Data Warehouse environment for these jobs. If you want to run only Hive ACID compaction tasks, you can choose to use either the Cloudera Data Warehouse or Cloudera Data Hub environments.

If you want to run the compaction jobs without changing the environment, it is recommended that you use Cloudera Data Warehouse. To avoid interference from Cloudera Data Hub, change the value of the hive.compactor.worker.threads Hive Server (HS2) property to '0'. This ensures that the compaction jobs are not processed by Cloudera Data Hub.

1. In Cloudera Manager, click Clusters Hive Configuration to navigate to the configuration page for HMS.
2. Search for hive.compactor.worker.threads and modify the value to '0'.
3. Save the changes and restart the Hive service.

DWX-17254: Merging Iceberg branches requires a target table alias

7.3.2

Hive supports only one level of qualifier when referencing columns. In other words only one dot is accepted. For example, select table.col from ...; is allowed. select db.table.col is not allowed. Using the merge statement to merge Iceberg branches without a target or source table alias causes an exception:

```
org.apache.hadoop.hive.ql.parse.SemanticException: ... Invalid t
able alias or column reference ...
```

Use an alias, for example t, for the target table.

```
merge into mydb.target.branch_branch1 t using mydb.source.branch
_branch1 s on t.id = s.id when matched then update set value = '
matched';
```

Apache Jira: HIVE-28055

DWX-17210, DWX-13733: Timeout issue querying Iceberg tables from Hive

7.3.2

When querying Iceberg tables from Hive, the queries can faile due to a timeout issue.

1. Add the following configurations to hadoop-core-site for the Database Catalog and the Virtual Warehouse.
   - fs.s3.maxConnections=1000
   - fs.s3a.connection.maximum=1000
2. Restart the Database Catalog and Virtual Warehouse.

DWX-14163: Limitations reading Iceberg tables in Avro file format from Impala

7.3.2

The Avro, Impala, and Iceberg specifications describe some limitations related to Avro, and those limitations exist in Cloudera. In addition to these, the DECIMAL type is not supported in this release.

None.

DEX-7946: Data loss during migration of a Hive table to Iceberg

7.3.2

In this release, by default the table property 'external.table.purge' is set to true, which deletes the table data and metadata if you drop the table during migration from Hive to Iceberg.

Either one of the following workarounds prevents data loss during table migration:

- Set the table property 'external.table.purge'='FALSE'.
- Do not drop a table during migration from Hive to Iceberg.

DWX-13062: Converting a Hive table having CHAR or VARCHAR columns to Iceberg causes an exception

7.3.2

CHAR and VARCHAR data can be shorter than the length specified by the data type. Remaining characters are padded with spaces. Data is converted to a string in Iceberg. This process can yield incorrect results when you query the converted Iceberg table.

Change columns from CHAR or VARCHAR to string types before converting the Hive table to Iceberg.

Apache Jira: HIVE-26507

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-92182: Inserting into Hive Iceberg tables on S3 fails with RazS3ClientCredentialsException

7.3.2, 7.3.1.706, 7.3.1.600, 7.3.1.500

In RAZ-enabled clusters where HDFS is the default file system, attempts to insert data into Hive Iceberg tables pointing explicitly to S3 locations fail with a RazS3ClientCredentialsException

None.

CDPD-89390/CDPD-83022: Incorrect row count displayed in table metadata after compaction

7.3.2, 7.3.1.600, 7.3.1.500, 7.3.1.400

After running data compaction operations on large tables, the row count displayed by the DESC RIBE FORMATTED command may be inaccurate. Initially, the count may appear higher than the actual number of rows. Subsequently, after running the ANALYZE command to update table statistics, the count might then appear lower than the actual number of rows.

This issue has been observed in large tables containing a significant number of historical snapshots (exceeding 10000). All these snapshots are primarily generated through UPDATE operations.

It is important to note that this is just a metadata display issue, and there is no loss of data. The underlying table data remains complete and correct.

To obtain an accurate row count, use the SELECT COUNT(*) query.

Known Issues in Iceberg REST Catalog

Known issues and technical limitations for Cloudera Iceberg REST Catalog are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

CDPD-123217: Hive Metastore TLS appears enabled in Cloudera Manager but the metastore log reports SSL disabled

7.3.2.20000

After you enable Auto TLS or configure TLS for the Hive Metastore service in Cloudera Manager, the UI can show Hive Metastore TLS as enabled even when the ENCRYPT_ALL_PORTS feature is not active on the Cloudera Manager server (the CMF_FF_ENCRYPT_ALL_PORTS feature flag). In that case the Hive Metastore process does not enable Thrift SSL, and the metastore log correctly reports Enable SSL = false at startup. The mismatch makes it difficult to confirm TLS when you validate on-premises Data Sharing and Cloudera Iceberg REST Catalog connectivity.

If you require Hive Metastore TLS, set export CMF_FF_ENCRYPT_ALL_PORTS=true and, when you use advanced TLS controls, export CMF_FF_TLS_ADVANCED_CONTROL=true in /etc/default/cloudera-scm-server on the Cloudera Manager server host. Restart the Cloudera Manager server, configure Hive Metastore TLS, redeploy the affected roles, and confirm the metastore log reports SSL enabled. Until a future Cloudera Manager release corrects the UI, treat the metastore log as authoritative when the encrypt-all-ports feature is disabled.

Known Issues in Apache Impala

Known issues and technical limitations for Impala are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

CDPD-121808: Concurrent GRANT statements issued through `impala-shell` can fail with raw database errors

7.3.2.20000

When multiple GRANT (or REVOKE) statements were issued concurrently through `impala-shell` against the same or related database objects, the policy creation logic in the authorization service was not thread-safe. Concurrent requests that attempted to create a new policy row at the same time could collide on a database uniqueness constraint. Instead of retrying or returning an actionable conflict response, the authorization service returned an HTTP 400 Bad Request error that exposed internal JDBC or persistence-layer stack trace details, rather than a clear, user-facing error message.

Avoid issuing concurrent GRANT or REVOKE statements against the same or related resources from multiple sessions at the same time. Issuing the same statements sequentially completed successfully.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

Impala Kerberos authentication failures in FIPS clusters

7.3.2.10000 SP1

Impala Kerberos authentication fails through SPNEGO after running for several hours in Red Hat Enterprise Linux (RHEL) 9 FIPS clusters. This behavior occurs because of a known Red Hat issue (RHEL-180783). Until Red Hat releases a permanent fix, if a user upgrades the krb5 OS packages to the latest version on RHEL 9 in FIPS mode, it may trigger this bug and break authentication.

Downgrade the krb5 OS packages to a stable version (such as 1.21.1-6) provided by Red Hat.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

CDPD-111251: `impala-shell` connections fail with a Thrift error over the HS2-HTTP protocol on SLES 15 with Python 3.11

7.3.2.100 CHF 1

On SLES 15, `impala-shell` connections that use the --protocol=hs2-http option fail when `impala-shell` runs with Python 3.11, with errors similar to the following:

```
  Num remaining tries: 3 PY_SSIZE_T_CLEAN macro must be defined
for '#' formats
```

This issue does not occur when `impala-shell` runs with the SLES 15 default Python interpreter, Python 3.6, or when you connect by using a protocol other than HS2-HTTP.

On SLES 15, run `impala-shell` with the default Python 3.6 interpreter instead of Python 3.11 when you connect by using the --protocol=hs2-http option.

Known issues identified in Cloudera Runtime 7.3.2

CDPD-90807: Thrift protocol limitation during Impala zero downtime upgrade (ZDU)

7.3.1.500 through 7.3.1.706, 7.3.2 and its higher versions

Zero Downtime Upgrades (ZDU) for Impala are not supported when upgrading to version 7.3.2. This is due to Thrift protocol incompatibilities that can cause queries to fail during upgrade.

None

CDPD-90250: Incorrect file storage location for Impala tables when using S3 as default filesystem

7.3.2

When you created an external table in Impala by using the LOCATION parameter, the files were stored in S3 instead of HDFS. This occurred even if the provided path was intended for HDFS, provided that S3 was configured as the default filesystem (fs.defaultFS) in the cluster. This inconsistency leads to AccessDeniedException errors if the Impala service does not have the necessary Ranger permissions to write to the S3 bucket.

None

IMPALA-14472: Writing arrays to Kudu tables

7.3.2

Impala does not currently support writing array data into Kudu tables.

None

Apache Jira: IMPALA-14472

Known issues identified before Cloudera Runtime 7.3.2

DWX-20490: Impala queries fail with "Caught exception The read operation timed out, type=<class 'socket.timeout'> in ExecuteStatement"

7.3.1.500

Queries in impala-shell fail with a socket timeout error in execute statement which submits the query to the coordinator. The error occurs when query execution takes longer to start mainly when query planning is slow due to frequent metadata changes.

Increase the socket timeout on the client side. Set --client_connect_timeout_ms to a higher value, e.g. add --client_connect_timeout_ms=600000 to the impala-shell command line.

DWX-20491: Impala queries fail with EOFException: End of file reached before reading fully

7.3.1.500

Impala queries fail with an EOFException when reading from an HDFS file stored in an S3A location. The error occurs when the file is removed. If the file is removed using SQL commands like `DROP PARTITION`, there may be a significant lag in Hive Metastore event processing. If removed by non-SQL operations, run `REFRESH` or `INVALIDATE METADATA` on the table to resolve the issue.

Run REFRESH/INVALIDATE METADATA <table>;

CDPD-94720: Impala startup failure due to invalid TLS v1.3 ciphers

7.3.1 and its higher versions

When running Impala on a machine with OpenSSL 1.1.1, providing an invalid or TLS v1.2 ciphersuite in the --tls_ciphersuites startup flag causes the process to fail during startup. While OpenSSL 3.x ignores invalid ciphers, OpenSSL 1.1.1 returns an error if any ciphersuite in the list is invalid, even if other valid TLS v1.3 ciphers are present.

Ensure that the list in the --tls_ciphersuites startup flag contains only valid TLS v1.3 ciphersuites and does not contain any TLS v1.2 ciphersuites.

Apache Jira: IMPALA-14625

IMPALA-532: Impala should tolerate bad locale settings

7.3.1 and its higher versions

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-691: Process mem limit does not account for the JVM's memory usage

7.3.1 and its higher versions

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-635: Avro Scanner fails to parse some schemas

7.3.1 and its higher versions

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-1024: Impala BE cannot parse Avro schema that contains a trailing semi-colon

7.3.1 and its higher versions

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

7.3.1 and its higher versions

When comparing a CHAR column value to a string literal, the literal value is not blank-padded and so the comparison might fail when it should match.

Use the RPAD() function to blank-pad literals compared with CHAR columns to the expected length.

IMPALA-1821: Casting scenarios with invalid/inconsistent results

7.3.1 and its higher versions

Using a CAST() function to convert large literal values to smaller types, or to convert special values such as NaN or Inf, produces values not consistent with other database systems. This could lead to unexpected results from queries.

None

IMPALA-2005: A failed CTAS does not drop the table if the insert fails

7.3.1 and its higher versions

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

7.3.1 and its higher versions

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

7.3.1 and its higher versions

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

7.3.1 and its higher versions

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

None

IMPALA-7072: Impala does not support Heimdal Kerberos

None

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.

IMPALA-2422: % escaping does not work correctly when occurs at the end in a LIKE clause

7.3.1 and its higher versions

If the final character in the RHS argument of a LIKE operator is an escaped \% character, it does not match a % final character of the LHS argument.

None

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

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

None

CDPD-59625: Impala shell in RHEL 9 with Python 2 as default does not work

7.1.9, 7.3.1 and its higher version

If you try to run impala-shell on RHEL 9 by setting the default python executable available in PATH to Python 2, it will fail since RHEL 9 is compatible only with Python 3.

If you run into such issues, set this parameter pointing to Python 3, IMPALA_PYTHON_EXECUTABLE=python3.

Impala cannot update table if the 'external.table.purge' property is not set to true

Impala cannot update a table using DDL statements if the 'external.table.purge' property is FALSE. ALTER TABLE statements return success with no changes to the table.

ALTER TABLE statements should be issued twice if "external.table.purge" was FALSE initially.



Note: For Iceberg tables, Impala users should always use CREATE TABLE since the Impala CREATE TABLE statement sets the flag to TRUE. However, Impala CREATE EXTERNAL TABLE sets the flag to FALSE.

Impala's known limitation when querying compacted tables

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

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

Use the REFRESH/INVALIDATE statements on the affected table to overcome the 'File does not exist' exception.

Impala api calls via knox require configuration if the knox customized kerberos principal name is a default service user name

To access impala api calls via knox, if the knox customized kerberos principal name is a default service user name, then configure "authorized_proxy_user_config" by clicking Clusters->impala->configuration. Include the knox customized kerberos principal name in the comma separated list of values <knox_custom_kerberos_principal_name>=*" where <knox_custom_kerberos_principal_name> is the value of the Kerberos Principal in the Knox service. Select Clusters>Knox>Configuration and search for Kerberos Principal to display this value.

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes.

7.1.7

You must use a single Knox node to access Impala UI.

CDPD-21828: Multiple permission assignment through grant is not working

7.1.7

None

IMPALA-11871: INSERT statement does not respect Ranger policies for HDFS

7.3.1, 7.3.1.300 and its higher version

7.3.1.100, 7.3.1.200

In a cluster with Ranger auth (and with legacy catalog mode), even if you provide RWX to cm_hdfs -> all-path for the user impala, inserting into a table whose HDFS POSIX permissions happen to exclude impala access will result in "AnalysisException: Unable to INSERT into target table (default.t1) because Impala does not have WRITE access to HDFS location: hdfs://XXXXXXXXXXX"

OPSAPS-46641: A single parameter exists in Cloudera Manager for specifying the Impala Daemon Load Balancer. Because BDR and Hue need to use different ports when connecting to the load balancer, it is not possible to configure the load balancer value so that BDR and Hue will work correctly in the same cluster.

The workaround is to use the load balancer configuration either without a port specification, or with the Beeswax port: this will configure BDR. To configure Hue use the "Hue Server Advanced Configuration Snippet (Safety Valve) for impalad_flags" to specify the the load balancer address with the HiveServer2 port.

Impala known limitation when querying compacted tables

7.3.1 and its higher versions

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.

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

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

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

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

Impala is not producing these logs.

None

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 SQLGetDa ta call for column 1 returns NULL.

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

Known Issues in Apache Kafka

Known issues and technical limitations for Kafka are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2

CDPD-107343: Changing the KRaft listener port with a normal restart breaks the Kafka cluster

7.3.2

When the KRaft listener port is changed for all KRaft roles simultaneously and the cluster is restarted using a normal restart, the Kafka cluster stops working. When all controllers change their listeners at the same time, they only know the old listeners of the other controllers. As a result, controllers cannot send metadata update requests or elect a new active controller, and the metadata topic is not updated. Brokers become unhealthy because they cannot reach any controller. This issue also affects enabling or disabling AutoTLS because that operation changes the listener port.

Use a rolling restart when changing the KRaft listener port. If a rolling restart is not an option, perform the change in two steps to avoid simultaneous listener changes across all controllers:

1. Add the new listener to the Kafka configuration alongside the existing listener.
2. Restart the cluster.
3. Remove the old listener from the Kafka configuration.
4. Restart the cluster.

You can add multiple listeners in the Kafka Broker Advanced Configuration Snippet in Cloudera Manager.

OPSAPS-79006: Kafka broker recommission command fails

7.3.2

7.3.2.10000 SP1

After you decommission a Kafka broker, the Recommission role action fails in Cloudera Manager. The command fails during host IP resolution because the required HOST environment variable is not set. As a result, you cannot return a decommissioned broker to the cluster using recommission in Cloudera Runtime 7.3.2.

Do not decommission a broker if you might need to recommission it later. If you must remove capacity, complete broker removal (decommission and delete) instead of relying on recommission.

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

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

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

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

Streams Messaging Manager cannot identify multiple listeners and still points to bootstrap server using the default broker port (9093 for SASL_SSL). You need to override the bootstrap server URL by performing the following steps:

1. In Cloudera Manager, go to Streams Messaging Manager Configuration Streams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve)

2. Override bootstrap server URL (hostname:port as set in the listeners for broker) for streams-messaging-manager.yaml.
3. Save your changes.
4. Restart SMM.

KAFKA-2561: Performance degradation when SSL Is enabled

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

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.

RANGER-3809: Idempotent Kafka producer fails to initialize due to an authorization failure

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

Kafka producers that have idempotence enabled require the Idempotent Write permission to be set on the cluster resource in Ranger. If permission is not given, the client fails to initialize and an error similar to the following is thrown:

```
org.apache.kafka.common.KafkaException: Cannot execute transacti
onal method because we are in an error state
              at org.apache.kafka.clients.producer.internals.Tr
ansactionManager.maybeFailWithError(TransactionManager.java:1125)
              at org.apache.kafka.clients.producer.internals.T
ransactionManager.maybeAddPartition(TransactionManager.java:442)
              at org.apache.kafka.clients.producer.KafkaProduce
r.doSend(KafkaProducer.java:1000)
              at org.apache.kafka.clients.producer.KafkaProduc
er.send(KafkaProducer.java:914)
              at org.apache.kafka.clients.producer.KafkaProducer
.send(KafkaProducer.java:800)
              .
              .
              .
              Caused by: org.apache.kafka.common.errors.Cluste
rAuthorizationException: Cluster authorization failed.
```

Idempotence is enabled by default for clients in Kafka 3.0.1, 3.1.1, and any version after 3.1.1. This means that any client updated to 3.0.1, 3.1.1, or any version after 3.1.1 is affected by this issue.

This issue has two workarounds, do either of the following:

- Explicitly disable idempotence for the producers. This can be done by setting enable.idempoten ce to false.
- Update your policies in Ranger and ensure that producers have Idempotent Write permission on the cluster resource.

CDPD-49304: AvroConverter does not support composite default values

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2.0

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

CFM-3532: The Stateless NiFi Source, Stateless NiFi Sink, and HDFS Stateless Sink connectors cannot use Snappy compression

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

This issue only affects Stateless NiFi Source and Sink connectors if the connector is running a dataflow that uses a processor that uses Hadoop libraries and is configured to use Snappy compression. The HDFS Stateless Sink connector is only affected if the Compression Codec or Compression Codec for Parquet properties are set to SNAPPY.

If you are affected by this issue, errors similar to the following will be present in the logs.

```
Failed to write to HDFS due to java.lang.UnsatisfiedLinkError: o
rg.apache.hadoop.util.NativeCodeLoader.buildSupportsSnappy()
```

```
Failed to write to HDFS due to java.lang.RuntimeException: nativ
e snappy library not available: this version of libhadoop was bu
ilt without snappy support.
```

Download and deploy missing libraries.



Important: Ensure that you complete steps 1-11 on all Kafka Connect hosts. Additionally, ensure that the advanced configuration snippet in step 12 is configured for all Kafka Connect role instances.

1. Create the /opt/nativelibs directory.

   ```
   mkdir /opt/nativelibs
   ```

2. Change the owner to kafka.

   ```
   chown kafka:kafka /opt/nativelibs
   ```

3. Locate the directory containing the Hadoop native libraries and copy its contents to the directory you created.

   ```
   cp /opt/cloudera/parcels/CDH/lib/hadoop/lib/native/* /opt/na
   tivelibs
   ```

4. Verify that libsnappy.so was copied to the directory you created.
5. Remove the following from /opt/nativelibs.

   ```
   libhadoop.a
                     libhadoop.so
                     libhadoop.so.1.0.0
   ```

6. Run the following command.

   ```
   hadoop version
   ```

   The command returns the Hadoop version running in the cluster. Note down the first three digits in the version.
7. Go to https://archive.apache.org/dist/hadoop/common/ and download the Hadoop version that matches the first three digits of the version running in the cluster.

   For example, if your Hadoop version is 3.1.1.7.1.9.0-296, then you need to download Hadoop 3.1.1.
8. Extract the downloaded archive.
9. Copy the following libraries from the downloaded archive to /opt/nativelibs on the cluster host.

   ```
   libhadoop.a
                     libhadoop.so.1.0.0
   ```

The libraries are located in hadoop-*[***VERSION***]*/lib/native.

10. Create a symlink named libhadoop.so and point it to /opt/nativelibs/libhadoop.so.1.0.0.

```
ln -s /opt/nativelibs/libhadoop.so.1.0.0 /opt/nativelibs/lib
hadoop.so
```

11. Change the owner of every entry within /opt/nativelibs to kafka.

```
chown -h kafka:kafka /opt/nativelibs/*
```

12. In Cloudera Manager, go to Kafka service Configuration .
13. Add the following key-value pair to Kafka Connect Environment Advanced Configuration Snippet (Safety Valve).
    - Key: LD_LIBRARY_PATH
    - Value: /opt/nativelibs
14. Click Save Changes.
15. Restart the Kafka service.

Unsupported Features

The following Kafka features are not supported in Cloudera:

- 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.

Limitations

Collection of partition level metrics may cause Cloudera Manager performance to degrade

If the Kafka service operates with a large number of partitions, collection of partition level metrics may cause Cloudera Manager 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 Streams Messaging Manager to monitor Kafka or Cruise Control for rebalancing Kafka partitions, be aware that both Streams Messaging Manager and Cruise Control rely on partition level metrics. If partition level metric collection is disabled, Streams Messaging Manager will not be able to display information about partitions. In addition, Cruise Control will not operate properly.

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

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

   The Kafka service name is the value of $SERVICENAME.
2. Turn off the collection of partition level metrics.
   a. Go to Hosts Hosts Configuration .
   b. Find and configure the Cloudera Manager Agent Monitoring Advanced Configuration Snippet (Safety Valve) configuration property.

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

```
[KAFKA_SERVICE_NAME]_feature_send_broker_topic_partition_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.

Known Issues in Apache Knox

Known issues and technical limitations for Apache Knox are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

CDPD-97987: AWS SDK for Java v2 log messages are missing from Knox IDBroker logs

7.3.2.20000

AWS SDK for Java v2 log messages are missing from Knox IDBroker logs when IDBroker performs AWS credential brokering. The AWS SDK bundle includes its own SLF4J implementation that overrides Knox's logging configuration, suppressing log output from AWS SDK classes even when custom loggers are configured.

None.

CDPD-124561: Apache Knox Gateway default login page flashes before modern-themed login page is displayed

7.3.2.20000

When users open an Apache Knox Gateway login page, the default login page may flash briefly before the modern-themed login page appears. User authentication is not affected by this issue.

None.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

CDPD-124438: Apache Knox Gateway UIs load the gear and pin icons from the internet, breaking icons in air-gapped environments

7.3.2.10000, 7.3.2.20000

In air-gapped networks and in environments that block outbound internet access from Knox UI client browsers, Apache Knox Gateway UIs try to load the gear and pin icons from the internet. The browser cannot download the icons, so they render as literal ligature text instead of glyphs. Affected UIs include the Knox home page, gateway admin UI, token generation UI, and token management UI. Knox webshell UI is not affected.

None.

Known Issues identified in Cloudera Runtime 7.3.2.100 CHF 1

CDPD-104007: Tokens on Knox Token Management page might intermittently fail to display

7.3.2.100

Tokens created under the Knox Token Management page might intermittently fail to display. The issue occurs randomly upon refreshing the page, where the underlying API call to retrieve user tokens is not triggered, causing the page to appear empty.

Refresh the page again until the tokens are displayed.

Known Issues identified in Cloudera Runtime 7.3.2

The following section lists the known issues identified in this release:

OPSAPS-76713: Knox crashes under high concurrent load when using PAM authentication on RHEL 9 or Rocky Linux 9

7.3.2

On RHEL 9 or Rocky Linux 9 versions 9.1 through 9.7, a concurrency bug in libeconf 0.4.1 causes libpam to crash under high concurrent load. This affects Knox when the cdp-proxy-api topology uses PAM authentication. When Knox handles multiple concurrent requests, libpam might crash, causing Knox service to crash and requests to the Knox cdp-proxy-api topology to fail.

Use AD/LDAP authentication instead of PAM authentication for the cdp-proxy-api topology. In Cloudera Manager, go to the Knox service, select Configuration , and configure the gateway_api_authentication_provider parameter to use AD/LDAP authentication. For configuration instructions, see Configure Apache Knox authentication for AD/LDAP.

CDPD-98063: Knox service discovery fails in FIPS-enabled clusters with JDK 17 on RHEL 9

7.3.2

7.3.2.10000

Knox Gateway cannot connect to Cloudera Manager during initial service discovery or periodic configuration analysis in FIPS-enabled clusters running JDK 17 on RHEL 9.

None.

Known Issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-76490: Ranger API bulk resource deletion fails when proxied through Knox

7.3.1 and its CHFs, 7.3.2

7.3.2.10000

Ranger API bulk resource deletion fails when the request is proxied through Knox. When Ranger sends a DELETE request with a body for bulk resource deletion, Knox does not forward the request body according to RFC 9110, causing the operation to fail.

None.

CDPD-76294: Knox service can not be started in a large size Private Cloud Base Cloudera Manager cluster

7.3.1.100 through 7.3.1.700, 7.3.2

For a large size Private Cloud Base Cloudera Manager cluster installed with Cloudera Runtime 7.3.1.0, you might face the problem that Knox can not be started with the following error message:

```
Wait Until Knox Gateway Can Serve Requests failed on Knox Gateway
```

Increase the Knox configuration parameter Knox Gateway Initial/Max Heapsize from 1 GiB to 2 GiB or 4 GiB, depending on the cluster size. Then save changes and run Restart Stale Services. After these steps, the Knox service can be started.

CDPD-71751: Creation of alias from the Cloudera Manager UI fails on FIPS

7.1.9 SP1 and its CHFs, 7.3.1 and its CHFs, 7.3.2

Users attempting to create aliases through the Cloudera Manager UI face issues in FIPS.

The alias(es) can be created using the Knox CLI:

1. ssh to Knox host.

2. Export these directories: export KNOX_GATEWAY_DATA_DIR="/var/lib/knox/gateway/data"; export KNOX_GATEWAY_CONF_DIR="/var/lib/knox/gateway/conf"
3. Set the FIPS-specific options for the Knox CLI:

```
export KNOX_CLI_MEM_OPTS="--add-exports=java.base/sun.securi
ty.provider=bctls --add-exports=java.base/sun.security.provi
der=com.safelogic.cryptocomply.fips.core --add-modules=com.s
afelogic.cryptocomply.fips.core --add-modules=bctls --module-
path=<BCTLS_JARS_DIR> -Dcom.safelogic.cryptocomply.fips.appro
ved_only=true"
```

   <BCTLS_JARS_DIR> is the directory containing the SafeLogic bctls and fips core jar files.
4. Run the following command to create the alias: /opt/cloudera/parcels/CDH/lib/knox/bin/knoxcli.sh create-alias <ALIAS_NAME> <ALIAS_VALUE>
5. Verify the addition using /opt/cloudera/parcels/CDH/lib/knox/bin/knoxcli.sh list-alias.

For HA deployments, users must do it on every Knox host (whereas the Save Alias command applies the change to all hosts automatically).

CDPD-71305: Concurrent impala shell connection failure

7.1.9 SP1 and its CHFs, 7.3.1 and its CHFs, 7.3.2

If a user makes a concurrent impala-shell connection through Knox, then the connection fails.

Use only one Knox role.

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

7.1.9, 7.2.18, 7.3.1 and its CHFs, 7.3.2

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

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

CDPD-28431: Intermittent errors could be potentially encountered when Impala UI is accessed from multiple Knox nodes

7.1.9, 7.3.1 and its CHFs, 7.3.2

You must use a single Knox node to access Impala UI.

CDPD-22785: Improvements and issues need to be addressed in convert-topology knox cli command

7.1.9, 7.3.1 and its CHFs, 7.3.2

None.

Knox issue with JDK version

7.1.9, 7.3.1 and its CHFs, 7.3.2

jdk-1.8.0_391 is not supported.

Cloudera recommends using Cloudera supported JDKs.

CDPD-84236: Token generated by one Knox host fails with Unknown token error on another Knox host in Data Engineering High Availability clusters

7.3.1.400 through 7.3.1.700, 7.3.2

In Data Engineering High Availability clusters, a token generated by one Knox host may fail with an Unknown token error when accessed through another Knox host. This issue occurs due to a race condition in the PostgreSQL database, which prevents one of the Knox instances from properly initializing its configured token state service.

Restart Knox on all hosts.

Known Issues in Apache Kudu

Known issues and technical limitations for Kudu are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified Before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

You get "The user 'kudu' is not part of group 'hive' on the following hosts: " warning by the Host Inspector

If you are using fine grained authorization for Kudu, and you are also using Kudu-HMS integration with HDFS-Sentry sync, then you may get the "The user 'kudu' is not part of group 'hive' on the following hosts: " warning while upgrading.

Run the following command on all the HMS servers:

```
usermod -aG hive kudu
```

Known Issues in Livy

Review the list of known issues in Livy in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

CDPD-102241: Missing Livy for Spark 3 logo in Knox UI

7.3.2.10000 SP1

The Knox Gateway UI does not display the Livy for Spark 3 service logo.

None.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

There are no known issues carried forward that affect Livy in Cloudera Runtime 7.3.2.

Known Issues in Navigator Encrypt

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

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known Issues in Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

Failed to start navencrypt-mount.service

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 and its SPs and CHFs

After installing NavEncrypt on SLES, if the comand "systemctl status navencrypt-mount" fails with the error: "Failed to start navencrypt-mount.service: Unit navencrypt-mount.service failed to load: No such file or directory", the systemd files need to be reloaded..

Run the command :

```
(sudo) systemctl daemon-reload
```

Known Issues in Oozie

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

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known Issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known Issues identifed before Cloudera Runtime 7.3.2

CDPD-93327: Passwords may be printed in the atlas-application.properties file

7.3.2

The /etc/atlas/conf/atlas-application.properties file is a publically readable file, and it may contain the Java trust store password. This file is read by the Oozie server and is propagated to the running actions. This allows the trust store password in the /etc/atlas/conf/atlas-application.properties file to be printed in the Oozie Launcher AM's (Yarn application) console logs, and the whole atlas-application.properties file is copied into the Oozie Launcher AM's (Yarn application) local directory.

None

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

7.3.2

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.

Workaround: 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.

Known Issues in Ozone

Known issues and technical limitations for Ozone are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

CDPD-105709: Recon UI displays inconsistent Solr health status

7.3.2.10000

The Recon Overview page incorrectly displays SOLR is HEALTHY as an alert when the service functions normally. Additionally, other pages within the Recon UI report SOLR is UNHE ALTHY despite Solr remaining in a healthy state.

None

Apache JIRA: HDDS-15571

CDPD-109005: Storage Container Manager (SCM) fails to update DatanodeDetails when only port configurations change

7.3.2.10000

When a DataNode re-registers with SCM or sends heartbeats after gaining new ports, SCM fails to update the stored DatanodeDetails in the node registry. SCM triggers an update only if the IP address, hostname, or software version changes. Consequently, SCM ignores port-only updates, which leads to stale metadata and potential failures in features that rely on the new ports.

None

CDPD-109002: Ozone DataStream writes fail after enabling Ratis DataStream on existing clusters

7.3.2.10000

When Ratis DataStream is enabled on an existing cluster, open pipelines retain stale DataNode details that lack the required DataStream port information. Consequently, Ozone client writes fail when they exceed the automatic threshold because the client cannot identify the necessary port for pipeline members. This occurs because SCM does not refresh member ports for existing open pipelines after the configuration change.

None

Apache JIRA: HDDS-15799

CDPD-99722: Container data checksum reverts to a stale value after successful reconciliation

7.3.2.0 and 7.3.2.10000

In Ozone, the BackgroundContainerDataScanner service reverts the container data checksum to a stale value even after a successful reconciliation repairs corrupt chunks. This occurs because the scanner builds a checksum tree in memory while reconciliation updates the container data on disk, causing the scanner to persist an outdated tree.

The issue resolves automatically during the next container scan cycle. To resolve the issue immediately, manually trigger container reconciliation from the command line.

Apache JIRA: HDDS-14936

CDPD-105115: DataNodes fail to communicate with SCM after restart

7.3.2.10000

In a SCM High Availability (HA) cluster, a restarting or failing-over follower SCM starts its DataNode protocol server before completing its Ratis log replay. This causes SCM to process DataNode container reports against an incomplete state, which leads to dropped replica locations. Consequently, if that follower becomes the leader, affected containers appear to have missing or empty replicas, and DataNode logs record errors such as Unable to clear disk, Failed to flush DB b efore close, and EOFExceptions.

This issue is temporary and typically resolves during the next report cycle when SCM rebuilds the replica list.

CDPD-119561: Ozone Manager (OM) becomes unresponsive during sustained snapshot operations

7.3.2.10000

The OM leader enters a deadlock state during heavy snapshot creation and deletion loads. This occurs when the Key Deleting Service holds a snapshot database lock while waiting for a transaction that cannot complete because the flush thread is blocked. As a result, the OM stops responding to client requests and snapshot operations time out.

- Restart the affected OM to release the locks and restore service.
- Limit snapshot deletion and heavy write activity to reduce the risk of recurrence.
- Increase the ozone.om.unflushed.transaction.max.count configuration to lower the probability of the deadlock (monitor for increased memory usage).

Apache JIRA: HDDS-16122

CDPD-104229: Delay in deleting snapshot YAML file by OM Snapshot Local Data Manager

7.3.2.10000

A significant delay of approximately 20 to 30 minutes occurs in the deletion of snapshot YAML files by the OM Snapshot Local Data Manager after the system processes a snapshot deletion request. This delay happens because the current design requires the inDegree, which represents the number of descendant snapshots, to reach zero before the system removes the YAML file. This behavior impacts subsequent operations, such as bootstrap test cases, that rely on the timely cleanup of these files.

None

CDPD-119610: OM leader remains in LEADER_AND_NOT_READY state during Snapshot Defragmentation Bootstrap

7.3.2.10000

The Ozone Manager (OM) leader remains in the LEADER_AND_NOT_READY state after a follower node stops for a planned bootstrap. This issue occurs due to a deadlock between the snapshot installation thread and the double-buffer flush thread on the follower node. Consequently, the cluster loses its usable quorum, and all OM client RPC requests fail with an OMLeaderNotR eadyException.

Restart the affected OM follower node to break the deadlock and restore normal cluster operations.

CDPD-119172: OM follower node bootstrap fails when OM leader changes

7.3.2.10000

In an Ozone Manager (OM) High Availability (HA) setup, the bootstrap process for a follower node fails if the OM leader node changes while the bootstrap is in progress. If the original leader stops (for example, through SIGTERM), the bootstrap thread attempts to download the database snapshot from the former leader and fails with a Connection refused error. When the thread subsequently attempts to connect to the newly elected leader, it receives a 503 Service Unavailable error because the new leader is not yet ready to serve requests.

To ensure a successful bootstrap, verify that two of the three Ozone Managers (OMs) are running before bootstrapping the third node.

CDPD-102216: Toggling to Old UI from the Cluster Capacity tab results in a 404 error

7.3.2.10000

The Cluster Capacity feature is available only in the new Recon UI. If you view the Cluster Capacity page in the new UI and click Switch to Old UI, the application attempts to load the same path in the legacy interface. Because this path does not exist in the old UI, a 404 Page Not Found error appears.

Manually navigate to the main dashboard if you encounter a 404 Page Not Found error after switching to the old UI. Use the new UI to view cluster capacity information, as the legacy interface does not support this feature.

Apache JIRA: HDDS-15148

CDPD-106805: Requests to Ozone S3 Gateway through HAProxy fail with 400 Bad Request during directory creation

7.3.2.10000

When using the S3 connector to access Ozone OBS buckets through an S3 Gateway behind HAProxy, requests fail with a 400 Bad Request error. This issue occurs because zero-byte directory-marker PUT requests (which include Expect: 100-continue headers) cause the S3 Gateway to skip reading the request body and return 200 OK immediately. This leaves unconsumed bytes on the wire and desynchronizes keep-alive connections. HAProxy then misinterprets subsequent requests, causing a parsing error.

Disable keep-alive connections in the HAProxy configuration by setting option httpclose and increasing the maxconn value (for example, to 4096).

CDPD-107197: The `ozone admin datanode diskbalancer status` command fails when passing a DataNode ID instead of an address

7.3.2.10000

The `ozone admin datanode diskbalancer status` command accepts only a DataNode network address or host string. When you use a valid datanodeUuid (DataNode ID) to identify the node, the command fails with an UnknownHostException or an invalid argument error. This behavior is inconsistent with other `ozone admin` commands that support both UUIDs and addresses for node identification.

To check the disk balancer status, use the DataNode network address or hostname instead of its UUID. Run the `ozone admin datanode list` command to locate the network address.

Apache JIRA: HDDS-15690

CDPD-100716: Reading keys results in a NO_REPLICA_FOUND error after SCM follower restart and leader transfer

7.3.2.10000

When a Storage Container Manager (SCM) follower restarts, it accepts DataNode container reports before catching up with the Ratis log. Processing these reports against a stale database triggers a NotLeaderException, causing the SCM to skip updating in-memory container replica locations. If this follower is later promoted to leader, reading keys fails with a NO_REPLICA_FOUND error due to missing replica metadata.

None

Apache JIRA: HDDS-14989

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2

CDPD-95733: Snapshot pruning tests fail due to empty compaction SST backups in Ozone in Cloudera Runtime7.3.2

7.3.2

In Ozone in Cloudera Runtime 7.3.2, snapshot pruning tests fail because the Ozone Manager (OM) RocksDB Compaction Service configuration key requires a new name. The property key changes from ozone.compaction.service.enabled (used in version 7.1.9) to ozone.om.compaction.service.enabled in version 7.3.2. If you use the old configuration key, the compaction service remains disabled by default. This causes empty compaction SST backup directories and fails the pruning operations.

Update the configuration to use the new property key. Set ozone.om.compaction.service.enabled to true to enable the compaction service in version 7.3.2.

OPSAPS-76062: Exposing "ozone.replication" in Cloudera Manager configurations and assigning a default value

7.3.2

Exposing "ozone.replication in Cloudera Manager configurations and assigning a default value to it is causing it to override bucket replication configuration as a client side configuration even when you does not mean to set the client side configurations.

None

CDPD-73792: The new snapshot could not be found after renaming the old snapshot

7.3.2

The `ozone sh snapshot rename` command renames snapshots. It is a feature developed by the Apache Ozone community and is included in Cloudera Base on premises 7.3.1. However, it does not work properly, and Cloudera Base on premises does not support it.

None

Apache JIRA: HDDS-11384

CDPD-63350: Force deleting a FSO bucket and its contents while running `rb --force` from AWS S3 API is failing

7.3.2

Force deleting a File System Optimized (FSO) bucket and its contents while running the `rb --force` command from the AWS S3 API might fail with an error, because the S3 client sends individual delete requests for each key in the bucket. It might delete or fail to delete individual keys or directories, depending on the availability of leaf elements. It can completely delete the bucket

only when all the keys or directories are cleaned up. If some keys are not deleted, the bucket will not be deleted.

Sample error message

```
# aws s3 rm s3://buck-fso --recursive
delete: s3://buck-fso/dir1/
delete: s3://buck-fso/dir1/dir2/
delete: s3://buck-fso/dir3/dir4/dir5/
delete: s3://buck-fso/dir3/dir4/
delete failed: s3://buck-fso/dir3/
# Rerun the same command again
# aws s3 rm s3://buck-fso --recursive
delete: s3://buck-fso/dir1/
delete: s3://buck-fso/dir3/
# To Confirm
# ozone sh key list s3v/buck-fso
[ ]
```

Run the `rb --force` command multiple times to completely clean up the keys and directories.

Apache JIRA: HDDS-9637

CDPD-98751: Mismatched Replicas tab in the Recon UI fails to display containers with inconsistent replica checksums

7.3.2

In the Ozone Recon UI, the Mismatched Replicas tab does not update or display containers when one or more replicas have differing checksums. Instead, they are displayed in the Under-Replicated tab.

The replica can be checked in the Under-Replicated tab, or can be cross-checked against API response.

CDPD-97512: Mismatch in Open Key count between Overview page and OM DB Insight in the Recon UI

7.3.2

In the Ozone Recon UI, the Open Key count on the **Summary** section of the Overview page might not match the count on the OM DB Insight Open Key tab. Users might see different Open Key values for the same cluster across these two views.

Use the OM DB Insight Open Key tab for the most accurate Open Key count.

CDPD-97376: Container replication counts mismatch in Recon UI

7.3.2

In the Ozone Recon UI, container replication counts, including Under-Replicated, Over-Replicated, and Mis-Replicated counts, differ between the updated and the legacy **Container** page for the same cluster.

Refer to the legacy **Container** page to view accurate replication counts.

CDPD-97311: Incorrect Creation Time and Modification Time displayed on the Namespace Usage page in the Recon UI

7.3.2

In the Ozone Recon UI, the **Namespace Usage** page displays incorrect Creation Time and Modification Time timestamps. These values do not accurately reflect the actual creation or last modification times or dates of the namespace, resulting in inaccurate metadata information.

Retrieve the correct timestamp values directly from the API response.

CDPD-97312: Mismatch between cluster State Container count and Container Summary totals

7.3.2

The Ozone Recon UI can display discrepancies between the Storage Container Manager (SCM) container count and the Recon container summary. This occurs because Recon does not synchronize all container states such as QUASI_CLOSED leading to inconsistent totals.

No workaround within the Ozone Recon UI. Use `ozone admin container report` CLI command to obtain the correct container counts for all states in the Ozone cluster.

CDPD-99248: After cdh upgrade, Ozone encounters failure while running the `Finalize Upgrade for SCM on role Storage Container` command

7.3.2

When finalizing an Ozone upgrade for the first time from Cloudera Manager, the `Finalize Upgrade for SCM on role Storage Container` command might fail with the following stderr message:

```
"Invalid response from Storage Container Manager.
Current finalization status is: FINALIZATION_IN_PROGRESS"
```

This error occurs even though finalization continues to run on the Storage Container Manager (SCM).

Ignore the failure in Cloudera Manager. Use the `ozone admin scm finalizationstatus` command to monitor progress and wait for the process to complete on SCM.

CDPD-98892: File Size Distribution bucket size range calculation is not correct

7.3.2

If large number of buckets exists within a volume, the Ozone Recon UI might display incorrect file size distribution bucket size range calculation in the File Size Distribution chart on the **Insights** page.

None

Apache JIRA: HDDS-14827

CDPD-93116: Ozone client hangs intermittently when disks are full

7.3.2

This hang is caused by continuous write retries that persist until the pipeline on the Datanode closes.

The Ozone client hangs for approximately five minutes when writing data to a Datanode if a disk full exception or other Datanode error occurs. This hang is caused by continuous write retries that persist until the pipeline on the Datanode closes.

Control the request retry behavior by setting the following configurations on the client side:

Table 12: Client-side retry request configurations

Configuration	Recommended value
hdds.ratis.raft.client.rpc.request.timeout	30s
hdds.ratis.client.multilinear.random.retry.policy	1s, 1
hdds.ratis.client.exponential.backoff.max.sleep	5s
hdds.ratis.client.exponential.backoff.base.sleep	1s
hdds.ratis.client.exponential.backoff.max.retries	2

Apache JIRA: HDDS-14040

Known Issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-91562: test_validate_certs_configs configuration is failing with the maximum lifetime validation

7.3.2, 7.3.1.600

In daylight saving time zones, the autogenerated Ozone certificate duration might differ from the expected duration. This discrepancy is minor, because the default certificate duration is 365 days or five years, depending on the Ozone component.

CDPD-75954: The ozone debug `ldb` command and `ozone auditparser` fails with java.lang.UnsatisfiedLinkError

7.3.2, 7.3.1.400

For information on workaround, see Changing temporary path for Ozone services and CLI tools.

CDPD-54885: Ozone Prometheus does not work with TLS

7.3.2, 7.3.1 and its SPs and CHFs

The Prometheus service shipped by Ozone does not support TLS mode. So, Prometheus is not able to gather metrics from Ozone endpoints when TLS is enabled.

Go to Ozone Configuration Ozone Prometheus Endpoint Token and in the Ozone Prometheus Endpoint Token property enter any random string. This configuration generates a plaintext token in the Ozone endpoint process directory allowing Prometheus to authenticate and collect metrics despite the TLS limitation.

CDPD-56684: Keys and buckets get deleted without volume permission

7.3.2, 7.3.1 and its SPs and CHFs

When a volume deletion is initiated, the system recursively deletes all buckets and keys within the volume before attempting to delete the volume itself. Because the ACL check for volume deletion permissions occurs only in the end, all the data within the volume is deleted even without having delete permission on the volume.

Note: This occurs even if no specific bucket or key permissions were granted to the user for recursive deletion.

CDPD-50610: Large file uploads are slow with OPEN and stream data approach

7.3.2, 7.3.1 and its SPs and CHFs

Hue file browser uses the append operation for large files. This API is not supported by Ozone in 7.1.9, therefore large file uploads can be slow or can time out in the browser.

Use native Ozone client to upload large files instead of the Hue file browser.

OPSAPS-66469: Ozone-site.xml is missing if the host does not contain HDFS roles

7.3.2, 7.3.1 and its SPs and CHFs

The client side /etc/hadoop/conf/ozone-site.xml file is not generated by Cloudera Manager if the host does not have any HDFS role. Because of this, issuing Ozone commands from that host fails because it cannot find the service name to hostname mapping. When this issue occurs, an error message is displayed: # ozone sh volume list o3://ozoneabc 23/03/06 18:46:15 WARN ha.OMProxyInfo: OzoneManager address ozoneabc:9862 for serviceID null remains unresolved for node ID null Check your ozone-site.xml file to ensure ozone manager addresses are configured properly.

Add the HDFS gateway role on that host.

CDPD-63144: hadoop.ozone.om.request.key.OMKeyRenameRequestWithFSO: Rename key failed with "failed to get parent dir" error

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

Key rename inside the FSO bucket fails and discplays the Failed to get parent dir error. This happens when running impala workloads with ozone.

None.

CDPD-74331: Key put fails with "CompletionException: Failed to write chunk"

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

Key put fails and displays the Failed to write chunk error when there is a volume failure during configuration.

None.

CDPD-74475: YCSB test with Hbase on Ozone degrades performance

7.3.2, 7.3.1 and its SPs and CHFs

HBase on Ozone is currently provided as a Tech Preview feature. Performance characteristics, including throughput, may not yet match those of HBase running on HDFS.

None.

CDPD-74884: Exclusive size of snapshot is always returning 0

7.3.2, 7.3.1 and its SPs and CHFs

The exclusiveSize and exclusiveReplicationSize statistics provided by the snapshot info command return a value of 0 even when the snapshot contains exclusive keys or files that do not exist in other snapshots.

None.

Apache JIRA: HDDS-11528

CDPD-75042: AWS Cli recursive delete only deletes the leaf element

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

AWS CLI `rm` or `delete` command fails to delete all files and directories on the Ozone FSO bucket. It only deletes the leaf node.

None.

CDPD-75204: Setting file name length limit can cause NN shutdown with FSImage related error

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

Namenode restart fails after dfs.namenode.fs-limits.max-component-length is set to a lower value and there is existing data present which exceeds the length limit.

Increase the value for the dfs.namenode.fs-limits.max-component-length parameter and restart the namenode.

CDPD-75635: Ozone write fails intermittently because SCM remains in safe mode.

7.3.2, 7.3.1 and its SPs and CHFs

Ozone write operations can fail intermittently after restarting the Storage Container Manager (SCM) leader node (or when stopping the OM leader and an SCM follower). This occurs because SCM may remain in safe mode after the restart, causing write/block allocation prechecks to fail (for example, SafeModePrecheck failed for allocateBlock).

Wait for SCM to exit safe mode automatically after it meets the required thresholds. Alternatively, manually force SCM to exit from safe mode using CLI options.

Known Issues in Apache Parquet

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

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues in this release.

Known Issues in Apache Phoenix

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

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

OMID configuration issue after upgrading to Cloudera Runtime 7.3.2.

7.3.2.10000

After upgrading from Cloudera Runtime 7.1.9 or 7.3.2.0 to 7.3.2.10000, the OMID service is configured incorrectly. If HBase is configured to use the unsecure ZooKeeper port (HBase ZooKeeper Secure Client Enabled is set to false), creating an OMID transactional table from phoe nix-sqlline fails because the required TLS configuration is not sent to ZooKeeper.

Use one of the following workarounds:

- **Workaround 1**: On the upgraded cluster, set the OMID configuration property Enable TLS/SSL for Omid tso server (ssl_enabled) to false. Restart OMID and all dependent services with Re-deploy client configuration enabled.
- **Workaround 2**: On the upgraded cluster, set the HBase configuration property HBase ZooKeeper Secure Client Enabled (hbase.zookeeper.property.client.secure) to true. Restart HBase and all dependent services with Re-deploy client configuration enabled.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

There are no known issues carried forward that affect Phoenix in Cloudera Runtime 7.3.2.

Known Issues in Apache Ranger

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

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2

CDPD-125287: Service startup fails when Ranger audit logging to Log4j is enabled

7.3.2, 7.3.2.10000, 7.3.2.20000

When the Ranger plugin is configured to log access audit events to Log4j by setting xasecure.audit.destination.log4j=true in ranger-<service>-audit.xml, the service fails to start because the Log4JAuditDestination class cannot be found. The following error is logged:

```
java.lang.RuntimeException: Failed to create AuditDestination for
 class: org.apache.ranger.audit.destination.Log4JAuditDestination
Caused by: java.lang.ClassNotFoundException: org.apache.ranger.au
dit.destination.Log4JAuditDestination
```

Manually add the missing Ranger audit JAR to the affected service's Ranger plugin directory:

1. Obtain ranger-plugins-audit-2.6.0.7.3.2.0-957.jar from the CFM parcel.
2. Copy the JAR to the following directory:

   ```
   /opt/cloudera/parcels/CDH/lib/<service>/libs/ranger-<service>-
   plugin-impl/
   ```

3. Replace <service> with the name of the affected service, such as hive, hdfs, kafka, hbase, or solr.

CDPD-97234: Ranger RAZ client performance tuning configurations do not take effect

7.3.2, 7.3.2.10000

In Cloudera Runtime 7.3.2.0, the Ranger RAZ client performance tuning parameters documented for adjusting connection timeout and retry behavior may not take effect as expected. During an Availability Zone (AZ) failover, HDFS operations continue to use default timeout and retry values even after explicitly lowering them through the documented configuration. This occurs because some configuration properties, such as ranger.raz.client.max.retry, do not honor the expected configuration prefix, causing the custom values to be ignored.

To apply RAZ client performance tuning correctly, use the following configuration properties in your core-site settings:

```
<property>
    <name>ranger.raz.client.max.retry</name>
    <value>1</value>
</property>
<property>
    <name>fs.s3a.ext.raz.rest.client.connection.pool.retry-count
</name>
    <value>1</value>
</property>
<property>
    <name>fs.s3a.ext.raz.rest.client.connection.timeoutMs</name>
    <value>30000</value>
```

```
</property>
<property>
    <name>fs.s3a.ext.raz.rest.client.read.timeoutMs</name>
    <value>15000</value>
</property>
```

CDPD-101412: Ranger upgrade may fail in MySQL 8.x when the default charset is set to utf8mb4

7.3.2

In some environments, upgrading Ranger that uses an external MySQL 8.x database can fail if the server#level default character set is configured as utf8mb4. The upgrade may stop with database or schema#related errors because Ranger may not know the default charset of the MySQL server.

Create the required index manually by executing the following SQL statement in the Ranger DB:

```
CREATE INDEX x_trx_log_IDX_trx_id ON x_trx_log (trx_id(190));
```

This prefix index ensures the key length stays within the 3072-byte limit for utf8mb4 environments. For more details, check Additional Steps for Apache Ranger.

Known Issues identified before Cloudera Runtime 7.3.2

CDPD-106623: Ranger HDFS plugin incorrectly allows move-to-trash operations

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs

When ranger.hdfs.authz.enable.optimization is set to true (the default), the directory-rename optimization in the Ranger HDFS plugin incorrectly short-circuits permission evaluation to ALLOW during move-to-trash operations. As a result, users who have permissions on their own .Trash directory may be able to move or delete files and directories to trash even if they do not have the required permissions on the source directory.

Set the ranger.hdfs.authz.enable.optimization property to false in the Ranger HDFS plugin configuration (ranger-hdfs-security.xml). This disables the optimization logic and enforces normal permission checks for all operations, including move-to-trash.

CDPD-94671: Knox-proxied Ranger Swagger UI generates incorrect API URLs

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 and its SPs and CHFs

When you access the Ranger API Documentation (Swagger UI) through the Knox proxy (Knox UI Ranger UI API Documentation Swagger UI), the generated REST API URLs and curl commands use an incorrect endpoint format. Executing these API calls results in a 404 Undocumented Error: Not Found. This issue also occurs when running the Swagger-generated curl commands from the command line.

APIs accessed directly using the Ranger Admin host and port (rangeradminhost:6182) work as expected.

OPSAPS-75673: Wrong enablement of Ranger RMS Database Full Sync command

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs, 7.1.9 and its SPs and CHFs, 7.1.8 and its SPs and CHFs

The `Ranger RMS Database Full Sync` command should be enabled only when all RMS server instances are stopped. This is required to ensure that the RMS database synchronizes correctly without introducing conflicts or data corruption. However, when HA (High Availability) is enabled on the cluster, the command becomes available from Cloudera Manager Ranger RMS Actions drop-down, even though only one Ranger RMS instance is stopped while the others are still running.

None.

Ranger DB upgrade failure with password authentication error

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

During a cluster upgrade (such as upgrading to higher version), the Ranger PostgreSQL database upgrade step fails with a password authentication error similar to the following:

```
SQLException : SQL state: 28P01 org.postgresql.util.PSQLExceptio
n: FATAL: password authentication failed for user "ranger" Error
Code: 0
```

This issue occurs even if the password configured in Cloudera Manager matches the database. During the upgrade process, the installer executes a Jisql command where the password string is encapsulated in single quotes ('...'). If the password itself contains a single quote character, authentication fails.

To resolve this issue, remove single quote characters from the Ranger database user password and proceed with the upgrade:

1. Update the password in the database:

   Log in to your PostgreSQL/database host and alter the ranger user's password so that it does not contain single quotes (').
2. Update Cloudera Manager configuration:
   a. Go to the Ranger service in Cloudera Manager.
   b. Navigate to the Configuration tab.
   c. Search for Ranger Database User Password and enter the new password without single quotes.
   d. Save the configuration changes.
3. Resume the upgrade:

   Resume or retry the Ranger database upgrade process from Cloudera Manager.

CDPD-68806: The Revoke operation for users belonging to a group or role permission does not function as expected

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

List command is listing all the tables even when the user permission is revoked. And also the command does not add any deny policy to Ranger for that specific user.

This behavior is currently not supported in HBase shell. Must be handled manually using the Ranger policy change.

CDPD-68739: The revoke command does not work when using the HBase shell

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

While using the HBase shell, running the revoke command does not cancel the user permission. Users are able to perform actions even after running the revoke command.

None.

CDPD-58704: hadoop roll key/key delete command shows operation failed error when one KMS host is down, even when operation succeeds

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs, 7.1.9 and its SPs and CHFs

In case of rollover/delete, client sends one more (last after delete request) request to KMS instances to clean their cache and that too to all registered kms instances. if one KMS instance is stopped (not deleted), the client gets a runtime exception.

This simply returns the runtime exception on client end for stopped instances but doesn't break any functionality.

CDPD-56803: When there is no existing policy for user and a revoke request comes from hbase, then will get this error

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs, 7.1.9 and its SPs and CHFs

None.

CDPD-56741: Improvement in log message when jwtauth not used

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs, 7.1.9 and its SPs and CHFs

None.

CDPD-56738: Ranger RMS showing FileNotFoundException: /usr/share/java/oraclepki.jar in Oracle 19 setup

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs, 7.1.9 and its SPs and CHFs

This is a warning log printed in catalina.out file when Ranger RMS server is initialized. The following exception is observed only in Oracle 19 setup: FileNotFoundException: /usr/share/java/oraclepki.jar

None.

CDPD-48975: Ranger KMS KTS to KMS DB migration: keys with the same name but different case are not migrated

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs, 7.1.9 and its SPs and CHFs

KMS keys are not case sensitive.

No workaround. Such key combinations are very rare and the migration doc was updated to check such keys before starting the migration.

CDPD-42598: Kafka policy creation allowed with incorrect permissions

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs, 7.1.9 and its SPs and CHFs, 7.1.8 and its SPs and CHFs

When creating a Kafka policy from the UI, the permissions "Idempotent write"and "Cluster action" are not displayed as they are not applicable for the "topic" resource, but when creating a policy for the "topic" resource with the permissions "Idempotent write" and "Cluster Action", the policy is created successfully when the expected behaviour is that the policy creation must fail as the permission is not applicable for the Kafka topic resource.

None.

CDPD-41582: Atlas Resource Lookup : Classification for "entity-type" lists only classification for the following payload: {"resourceName": "classification", "userInput": "", "resources": {"classification": []}}

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs, 7.1.9 and its SPs and CHFs, 7.1.8 and its SPs and CHFs

Expectation is to return all the classifications. But the response has only "classification". Happens similarly for entity-label, entity-business-metadata.

None.

Known Issues in Ranger KMS

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

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

There are no known issues carried forward that affect Ranger KMS in Cloudera Runtime 7.3.2.

Known Issues in Schema Registry

Known issues and technical limitations for Schema Registry are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-40380: Authorization checking issue when Kerberos is disabled

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

Due to an issue in Ranger, when Kerberos is disabled then it is not possible to check authorization.

1. Open Schema Registry configuration in Cloudera Manager.
2. Find the ranger.plugin.schema-registry.service.name field.
3. Replace GENERATED_RANGER_SERVICE_NAME with the actual name of the service.
4. Restart the Schema Registry service.

CDPD-49304: AvroConverter does not support composite default values

7.1.7 and all its SP and CHF releases, 7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

AvroConverter cannot handle schemas containing a STRUCT type default value.

None.

OPSAPS-70971: Schema Registry does not have permissions to use Atlas after an upgrade

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

Following an upgrade, Schema Registry might not have the required permissions in Ranger to access Atlas. As a result, Schema Registry's integration with Atlas might not function in secure clusters where Ranger authorization is enabled.

1. Access the Ranger Console (Ranger Admin web UI).
2. Click the cm_atlas resource-based service.

3. Add the schemaregistry user to the all - * policies.
4. Click Manage Service Edit Service .
5. Add the schemaregistry user to the default.policy.users property.

OPSAPS-69317: Kafka Connect Rolling Restart Check fails if SSL Client authentication is required

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

The rolling restart action does not work in Kafka Connect when the ssl.client.auth option is set to required. The health check fails with a timeout which blocks restarting the subsequent Kafka Connect instances.

You can set ssl.client.auth to requested instead of required and initiate a rolling restart again. Alternatively, you can perform the rolling restart manually by restarting the Kafka Connect instances one-by-one and checking periodically whether the service endpoint is available before starting the next one.

Known Issues in Apache Solr

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

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

Creating Solr collection snapshots using solrctl fails on TLS-enabled clusters

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 and its SPs and CHFs

When you run the `solrctl` command to create a collection snapshot (such as solrctl collecti on --create-snapshot) on a cluster with TLS and Kerberos enabled, the command fails with an IOException occurred when talking to server at: error message.

This issue occurs because the TLS truststore location and password are not exported in the shell environment, preventing the solrctl client from establishing a secure connection to the Apache Solr server over HTTPS. Using the Solr REST API directly to create or list snapshots is not affected.

To resolve this issue, export the truststore path and password in the *JAVA_TOOL_OPTIONS* environment variable before you run the `solrctl` command:

```
JAVA_TOOL_OPTIONS
=
"-Djavax.net.ssl.trustStore=[***PATH_TO_TRUSTSTORE_FILE***] -Dja
vax.net.ssl.trustStorePassword=[***TRUSTSTORE_PASSWORD***]"
```

After setting the environment variable, run the `solrctl` command:

```
solrctl --jaas
```

```
[
***PATH_TO_JAAS_CONF***
]
collection --create-snapshot
[
***SNAPSHOT_NAME***
]
-c
[
***COLLECTION_NAME***
]
```

Known Issues in Spark

Review the list of known issues in Spark in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Known Issues Identified In Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues identified in this release.

Known Issues Identified In Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues identified in this release.

Known Issues Identified In Cloudera Runtime 7.3.2.100 CHF 1

CDPD-111272: MSCK REPAIR TABLE may perform redundant ALTER TABLE when partitions are already tracked

7.3.2.100 CHF 1

When you run MSCK REPAIR TABLE on a Hive table, Spark may issue a redundant ALTER TA BLE for partitions that are already tracked, which can increase repair time.

None.

Known issues identified in Cloudera Runtime 7.3.2

CDPD-67806: Migration of ORC table with timestamp column to Iceberg with spark.sql.timestampType=TIMESTAMP_NTZ fails

7.2.18, 7.3.1.x

The Hive External Catalog converts NTZ timestamp types back to TimeStampType, which Hive does not support, leading to migration failures.

Set spark.sql.iceberg.use-timestamp-without-timezone-in-new-tables and spark.sql.iceberg.handle -timestamp-without-timezone to true to migrate ORC tables to Iceberg with timestamp columns without timezone..

CDPD-90447: Table is stored in HDFS when cluster is enabled with raz and S3 is set to default storage system

7.3.2

Tables created on a RAZ-enabled clusters with S3 set as the default storage system are stored in HDFS instead of the specified S3 bucket.

None. S3 was removed as the default storage.

OPSAPS-75684: Spark fails due to Zookeeper Custom Kerberos Principal issue

7.3.2

Incorrect Zookeeper principal configuration and missing JVM property setup leads to SASL authentication failures.

When a custom Zookeeper principal is used, add the -Dzookeeper.sasl.client.username=[***USERNAME***] JVM argument to spark.*.defaultJavaOptions or spark.*.extraJavaOptions in spark-defaults.conf.

Known issues identified before Cloudera Runtime 7.3.2

CDPD-94393: RuntimeWarning Failed to add file" message appears even when Spark successfully loads files

7.3.2, 7.3.1.0 and higher

In both Spark 2 and 3, due to an exception when attempting to add files to the Python path, the RuntimeWarning: Failed to add file message appears even when the Python JAR file is successfully loaded.

None. You can safely ignore the message as the file is loaded successfully and the message does not affect job completion.

Spark 3: RAPIDS Accelerator is not available

7.3.2, 7.3.1.0 and higher

The RAPIDS Accelerator for Apache Spark is currently not available in Cloudera Runtime7.3.1

None.

The CHAR(n) type handled inconsistently, depending on whether the table is partitioned or not.

7.3.1

7.3.2, 7.3.1.100 CHF1 and higher

In upstream Spark 3 the spark.sql.legacy.charVarcharAsString configuration was introduced, but it does not solve all incompatibilities with Spark 2.

None. A new configuration spark.cloudera.legacy.charVarcharLegacyPadding will be introduced in a future version to keep compatibility with Spark 2, but it isn't available in 7.3.1.

Note: The CHAR type is legacy in SQL, and using it is discouraged. Cloudera recommends using VARCHAR or STRING instead.

Apache Jira: SPARK-33480

Known Issues in Spark Atlas Connector

Review the list of known issues in Spark Atlas Connector in Cloudera Runtime 7.3.2.0, its service packs and cumulative hotfixes.

Known Issues Identified In Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues identified in this release.

Known Issues Identified In Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues identified in this release.

Known Issues Identified In Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues identified in this release.

Known Issues Identified In Cloudera Runtime 7.3.2.0

There are no known issues identified in this release.

Known Issues Identified Before Cloudera Runtime 7.3.2.0

There are no new known issues identified in 7.3.1.x.

Known Issues for Sqoop

Known issues and technical limitations for Sqoop are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified Before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-54770: Unable to read Sqoop metastore created by an older HSQLDB version

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

If you have upgraded to Cloudera Base on premises 7.1.8 Cumulative hotfix 4 or higher versions, you may encounter issues in reading the Sqoop metastore that was created using an older version of HyperSQL Database (HSQLDB).

Cloudera upgraded the HSQLDB dependency from 1.8.0.10 to 2.7.1 and this causes incompatibility issues in Sqoop jobs that are stored in HSQLDB.

After upgrading to Cloudera Base on premises 7.1.8 Cumulative hotfix 4, you must upgrade the Sqoop metastore and convert the database files to a format that can easily be read by HSQLDB 2.7.1. For more information, see Troubleshooting Apache Sqoop issues.

CDPD-44431: Using direct mode causes problems

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

Using direct mode has several drawbacks:

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

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

Sqoop direct mode is disabled by default. However, if you still want to use it, enable it by either setting the sqoop.enable.deprecated.direct property globally in Cloudera Manager for Sqoop or by specifying it in the command-line through -Dsqoop.enable.deprecated.direct=true.

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

7.3.2 and its SPs and CHFs, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs, 7.1.7 SP1 and its CHFs

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

Parquet columns inadvertently renamed

Problem: 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

Problem: 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.

PARQUET-99

Known issues in Streams Messaging Manager

Known issues and technical limitations for Streams Messaging Manager are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known Issues identified in Cloudera Runtime 7.3.2

Alert policies based on the LEADER ELECTION PER SEC attribute do not trigger

7.3.2

If the monitored Kafka cluster is running in KRaft mode, alert policies based on the **LEADER ELECTION PER SEC** alert policy attribute will not trigger because the underlying metrics are not available from Kafka. The attribute remains visible and selectable in Streams Messaging Manager even if the monitored cluster is running in KRaft mode.

None.

CDPD-99371: Prometheus metrics do not function properly with KRaft mode

7.3.2

If the Kafka cluster monitored by Streams Messaging Manager is running in KRaft mode and Streams Messaging Manager is configured to use Prometheus as its metrics store, alerts related to the active controller do not work as expected. In KRaft mode, the broker_activecontrollerc ount metric is reported by KRaft controllers rather than brokers. This change in metric reporting can cause aggregation issues and lead to alerts triggering indefinitely. Additionally, adding the /api/prometheus-metrics endpoint from KRaft hosts to the Prometheus configuration can cause anomalies in other alerts.

Use Cloudera Manager as the metrics store for KRaft-based Kafka clusters. Alternatively, if using Prometheus, disable all alerts based on the broker_activecontrollercount metric.

7.3.2.10000 SP1

CDPD-96648: zstd compression type is missing from the compression.type drop-down menu on the Edit Topic form

7.3.2

When you edit a topic in the Streams Messaging Manager UI, zstd is not available in the compression.type drop-down menu on the Edit Topic form, even though Kafka supports zstd compression.

Change the topic compression type using the `kafka-topics` CLI tool.

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

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

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

Cloudera Manager does not display a Log Files menu for Streams Messaging Manager UI role (and Streams Messaging Manager UI logs cannot be displayed in the Cloudera Manager UI) because the logging type used by Streams Messaging Manager UI is not supported by Cloudera Manager.

View the Streams Messaging Manager UI logs on the host.

CDPD-39313: Some numbers are not rendered properly in Streams Messaging Manager UI

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

Very large numbers can be imprecisely represented on the UI. For example, bytes larger than 8 petabytes would lose precision.

None.

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

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

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

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

1. In Cloudera Manager, go to SMMConfigurationStreams Messaging Manager Rest Admin Server Advanced Configuration Snippet (Safety Valve) for streams-messaging-manager.yaml.
2. Add the following value for bootstrap servers.

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

3. Save your changes.
4. Restart Streams Messaging Manager.

CDPD-82560: Streams Messaging Manager known unprotected endpoint

7.1.7 and its SP and CHF releases, 7.1.9 and its SP and CHF releases, 7.3.1 and its SP and CHF releases, 7.3.2

The Streams Messaging Manager Frontend exposes the /cm-configs and /configs public endpoints without requiring authentication. These endpoints do not share any sensitive information.

None.

Known Issues in Streams Replication Manager

Known issues and technical limitations for Streams Replication Manager are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

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

7.1.7 and all its SP and CHF releases, 7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

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

7.1.7 and all its SP and CHF releases, 7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

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 Streams Messaging Manager 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: Streams Replication Manager may automatically re-create deleted topics on target clusters

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

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 Streams Replication Manager is no longer replicating the topic.
3. Delete the remote topic in the target cluster.

CDPD-80872: Streams Replication Manager replication-records-lag is incorrect on empty partitions with non-zero end offset

7.1.9 and all its SP and CHF releases, 7.3.1 and all its SP and CHF releases, 7.3.2

When a replicated partition is empty and its end offset is non-zero, the replication-records-lag metric is incorrectly reported as the end offset (instead of 0 or NaN).

Wait for a new message to be written into the partition. When it gets replicated, the metric is reported correctly.

Limitations

Streams Replication Manager cannot replicate Ranger authorization policies to or from Kafka clusters

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

Known Issues in YARN, YARN Queue Manager and MapReduce

Known issues and technical limitations for YARN and YARN Queue Manager are addressed in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

Known issues identified before Cloudera Runtime 7.3.2 include only unresolved issues from previous releases that continue to affect the Cloudera Runtime 7.3.2 base release.

CDPD-75652: Reverse DNS lookup fails for YARN but works for HDFS

7.3.2, 7.3.1 and its SPs and CHFs

When submitting a YARN application from a host without the correct DNS configuration (reverse DNS does not work for the YARN ResourceManager's host) Kerberos principal validation error, such as Server has invalid Kerberos principal. This is because the reverse DNS configuration does not work for YARN ResourceManager's host.

Add the following to YARN Service Advanced Configuration Snippet (Safety Valve) for yarn-site.xml:

```
<property>
```

```
<name>yarn.resourcemanager.principal.pattern</name>
<value>*</value>
</property>
```

COMPX-14682: Fix health check after Queue Manager restart

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

In some cases, after the QM webapp and config-service are restarted Cloudera Manager reports a healthy status after some time. However, the configuration service might not be ready and usable yet when this status is reported. Therefore, QM (webapp) is not ready and usable yet.

Wait for a minute or two after restart and then try QM.

COMPX-14820: Delete Queue and its Children throws "Queue capacity was reduced to zero, but failed to delete queue."

7.3.2, 7.3.1 and its SPs and CHFs, 7.1.9 SP1 and its CHFs

When trying to perform the operation "Delete Queue and its Children" on a queue that has one or more siblings, the operation fails as YARN has some constraints. If the queue performing the operation "Delete Queue and its Children" is a leaf node, then the operations succeeds.

None.

A fresh install of 7.3.1 or its SPs or CHFs does not allow user to bypass the Setup Database screen for YARN Queue Manager

7.3.2, 7.3.1 and its SPs and CHFs

YARN Queue Manager in Cloudera Base on premises 7.3.1 does not require you to install a PostGres database, therefore users should not see the Setup Database screen and should be able to skip the Setup Database screen. With this known issue, users who are conducting a fresh install of 7.3.1 or its SPs or CHFs are not able to bypass the Setup Database screen as expected.

1. When conducting a fresh install of YARN Queue Manager in 7.3.1 or its SPs or CHFs, you must ensure that you have both Cloudera and Cloudera Manager upgraded to 7.3.1.
2. When you reach the Setup Database screen in the Cloudera Manager installation wizard for Queue Manager, enter any dummy values for the following fields:
   a. Database name: configstore
   b. Database Username: dbuser
   c. Database Password: dbpassword
   d. Database Hostname: localhost

   YARN Queue Manager will not connect to PostGres with the above details and will fall back to the embedded database.
3. Run the following script command in a browser console to enable the Continue button:

   document.querySelector('.btn.next').removeAttribute('disabled');
4. Click Continue and proceed with the YARN Queue Manager installation.
5. After installation is complete, SSH into the host that has Queue Manager installed, and run this command: sed -i 's/migrationCompleted=true/migrationCompleted=false/' /var/lib/hadoop-yarn/migration.properties

   Note: Enable Queue Manager in the YARN configurations, and restart YARN.
6. Restart YARN Queue Manager.

20202 Database migration after enabling opt-in migration

7.3.2, 7.3.1 and its SPs and CHFs

When migrating from an H2 database to a PostgreSQL database in YARN Queue Manager after installation or upgrade, you might encounter an issue only when you have followed the following specific scenario:

- New install or upgrade to Cloudera 7.1.9, forcing migration from H2 to PostgreSQL database.
- Upgrade to Cloudera 7.1.9 CHF2, moving back to H2 database.
- Upgrade to Cloudera 7.1.9 SP1 with valid PostgreSQL connection details in Queue Manager configurations.

To avoid any issues during the upgrade to version Cloudera 7.1.9 SP1, ensure that PostgreSQL connection details are removed from the YARN database configuration if you prefer to continue using the H2 database.

Queue Manager does not open on using a custom user with a default Kerberos principal

7.3.2, 7.3.1 and its SPs and CHFs

If a custom user is used with the default Kerberos principal, the Queue Manager web UI displays an HTTP ERROR 400 error.

Ensure that the Queue Manager process_username property matches the YARN process_username property.

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

7.3.2, 7.3.1 and its SPs and CHFs

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

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

JobHistory URL mismatch after server relocation

7.3.2, 7.3.1 and its SPs and CHFs

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

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

YARN cannot start if Kerberos principal name is changed

7.3.2, 7.3.1 and its SPs and CHFs

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

There are two possible workarounds:

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

YARN Resource Manager UI does not display logs

7.3.2, 7.3.1 and its SPs and CHFs

When the parameter mapreduce.cluster.acl.enabled is set to true, the Yarn RM UI does display and logs and the Logs are not available message is displayed.

Set mapreduce.cluster.acl.enabled to false.

Known Issues in Apache ZooKeeper

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

Known issues identified in Cloudera Runtime 7.3.2.20000 SP2

There are no new known issues identified in this release.

Known issues identified in Cloudera Runtime 7.3.2.10000 SP1

CDPD-121220: OMID TSO fails to start when ZooKeeper plaintext client port 2181 is disabled

7.3.2.10000

When ZooKeeper plaintext client port 2181 is disabled by setting zookeeper_enable_client_port to false in secure/TLS-enabled environments, the OMID TSO (Transaction Status Oracle) role fails to start if Omid ZooKeeper Secure Client Enabled omid_enable_secure_zookeeper is set to false. OMID TSO attempts to connect to ZooKeeper using port 2181, resulting in a connection refused error Can't contact ZK cluster and prevents the service from initializing, thereby causing deployment and cluster startup health checks to time out.

Configure OMID to use the ZooKeeper secure port:

1. In Cloudera Manager, navigate to the OMID service.
2. Go to the Configuration tab.
3. Search for Omid ZooKeeper Secure Client Enabled (omid_enable_secure_zookeeper).
4. Select or set this property to true to allow OMID to connect using ZooKeeper port 2182.
5. Save changes and restart the OMID service, or restart stale services in the cluster.

Known issues identified in Cloudera Runtime 7.3.2.100 CHF 1

There are no new known issues in this release.

Known issues identified in Cloudera Runtime 7.3.2

There are no new known issues identified in this release.

Known issues identified before Cloudera Runtime 7.3.2

There are no new known issues from previous releases that affect ZooKeeper in Cloudera Runtime 7.3.2.

Behavioral Changes in Cloudera Runtime 7.3.2

Behavioral changes denote a marked change in behavior from the previously released version to this version of Cloudera Runtime.

Behavioral Changes in Atlas

Functional adjustments and behavioral updates for Atlas are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Atlas, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see Behavioral Changes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

Summary: Faster Atlas entity detail for large relationship graphs

Previous behavior:

When an entity had a very large relationship graph (such as some Apache Ozone buckets), the API returned all related entities in one response. This caused slow load times or an unresponsive Apache Atlas user interface, and the entity detail page did not display Properties, Classifications, Relationships, and Audits.

New behavior:

To optimize the Entity Detail page load time, the initial response now omits relationship attributes and referred entities. Details for the Schema tab, Relationships tab, and Terms are fetched on demand using a dedicated Relationship API. Key UI enhancements include:

- Relationships tab: The table view has been replaced with a card view that supports on-scroll pagination.
- Properties tab: Relationship details have been removed from the Properties tab and consolidated within the Relationships tab.

Cloudera Runtime 7.3.2.100 CHF 1

Summary: Faster Atlas entity detail for large relationship graphs

Previous behavior:

When an entity had a very large relationship graph (such as some Apache Ozone buckets), the API returned all related entities in one response. This caused slow load times or an unresponsive Apache Atlas user interface, and the entity detail page did not display Properties, Classifications, Relationships, and Audits.

New behavior:

To optimize the Entity Detail page load time, the initial response now omits relationship attributes and referred entities. Details for the Schema tab, Relationships tab, and Terms are fetched on demand using a dedicated Relationship API. Key UI enhancements include:

- Relationships tab: The table view has been replaced with a card view that supports on-scroll pagination.
- Properties tab: Relationship details have been removed from the Properties tab and consolidated within the Relationships tab.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Atlas, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see Behavioral Changes.

Summary: Automatic purging of soft-deleted entities is introduced

Previous behavior:

Previously, only the API call `DELETE /api/atlas/admin/purge` was available to manually purge soft-deleted entities. Additionally, `the DELETE api/atlas/v2/entity/guid/{{guid}}` API call could not delete the column lineages entities of Hive, Impala and Spark process entities. This could lead to sparse graphs resulting in reduced query performance.

New behavior:

A built-in auto-purge mechanism is introduced, deleted entities are purged in two stages. The first stage is a soft-delete at each Atlas startup. In the second-stage, soft-deleted process entities are purged based on a cron job. For more information, see Atlas Auto-Purging overview.

Summary:

Entity attributes details and sparkPlanDescription are no longer sent in the Spark process entity

Previous behavior:

The spark_process entity attributes details and sparkPlanDescription are populated with query plan details, which can contain a large amount of text, often in megabytes. This amount of data can incur unnecessary processing costs.

New behavior:

The atlas.spark.plan.enabled is set to false by default. Set it to true to send the details and sparkPlanDescription attributes in the Spark process entity. When these attributes are not sent, the cost of having large amount of data processed in Atlas is avoided.

Behavioral Changes in Avro

Functional adjustments and behavioral updates for Avro are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

There are no behavioral changes in this release.

Behavioral Changes in Cloud Connectors

Behavioral changes denote a marked change in behavior from the previously released version to this version of Cloud Connectors.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

There are no behavioral changes in this release.

Behavioral Changes in Cruise Control

Functional adjustments and behavioral updates for Cruise Control are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

Summary: ZooKeeper service dependency removed from Cruise Control

Previous behavior:

Cruise Control had a ZooKeeper service dependency in Cloudera Manager.

New behavior:

The ZooKeeper service dependency is removed.

Summary: Default JMX settings changed to restrict connections to localhost

Previous behavior:

No dedicated Cloudera Manager configuration option existed to control JMX JVM flags for Cruise Control. JMX connections were unrestricted by default.

New behavior:

A new CC_JMX_OPTS Cloudera Manager configuration option is available. Its default value restricts JMX connections to localhost only and enables SSL. If you previously set CC_JMX_OPTS, your custom value is preserved on upgrade. To revert to open JMX, update CC_JMX_OPTS in Cloudera Manager.

Summary: Default Supported Goals and Anomaly Detection Goals updated in Cloudera Manager

Previous behavior:

The Cloudera Manager defaults for Supported Goals did not include BrokerSetAwareGoal. The defaults for Anomaly Detection Goals covered rack, replica, and disk capacity only, not network inbound, network outbound, or CPU capacity goals.

New behavior:

The default Supported Goals list includes BrokerSetAwareGoal. The default Anomaly Detection Goals list also includes NetworkInboundCapacityGoal, NetworkOutboundCapacityGoal, and CpuCapacityGoal. Cloudera Manager defaults for these goal lists now match the Cruise Control defaults. No action is required after upgrade.

Behavioral Changes in HBase

Functional adjustments and behavioral updates for HBase are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

Summary: HBase enforces FIPS 140-3 compliant cryptographic standards

Previous behavior:

HBase used the default Java Cryptography Extension (JCE) provider, which permitted non-FIPS-compliant cryptographic algorithms. Deployments in FIPS-restricted environments had no platform-level enforcement, requiring manual configuration to restrict algorithm usage. TLS cipher suites and authentication token hashing were not validated against FIPS 140-3 requirements, which could result in failures during cryptosense compliance audits.

New behavior:

HBase is updated to use only FIPS 140-3 approved cryptographic algorithms. When the cluster operates in FIPS mode, non-compliant algorithms are disabled across all HBase services, including the HBase Controller (formerly known as HBase Master) and RegionServers. Only FIPS 140-3 approved TLS cipher suites are accepted for client connections and inter-service communication; non-compliant cipher suites are automatically excluded.

In version 7.3.2.10000, you might observe substantial performance degradation during HBase scan workloads on FIPS-enabled clusters utilizing AES encryption.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for HBase, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.700. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see Behavioral Changes.

Summary: HBOSS related configuration items are removed.

Previous behavior:

HBase contains the following HBOSS related configuration items.

- fs.hboss.fs.s3a.impl
- fs.hboss.sync.impl

New behavior:

The above configuration items are removed because HBOSS support is deprecated.

Summary: Removed the hbase.secure.rpc.engine configuration property

Previous behavior:

The hbase.secure.rpc.engine configuration property is obsolete because HBase performs replication securely by default.

New behavior:

The hbase.secure.rpc.engine configuration property is removed because it is obsolete and no longer needed.

Summary: The default values for the following configuration items are updated

Previous behavior:

Parameter name	Description	Default value
hbase.bucketcache.minfactor	The minimum percentage of BucketCache usage that should be targeted by the mass eviction process.	-
hbase.server.netty.tls.client.auth.mode	The client authentication mode for mTLS authentication when TLS RPC is in use. Three modes are NEED, WANT and NONE.	NEED
ip_version	Specifies the IP version the service must use for network communication. • IPv4 - Uses IPv4 exclusively. • IPv6 - Uses IPv6 exclusively. • Dual-stack (IPv4 &IPv6) - Supports both IPv4 and IPv6, enabling communication over both protocols.	-
hbase.client.netty.tls.enabled	Encrypt communication between clients and HBase RPC client mode using Transport Layer Security (TLS) (formerly known as Secure Socket Layer (SSL)).	-
hbase.bucketcache.acceptfactor	The percentage of BucketCache usage that must trigger the eviction of blocks.	-
hbase.rpc.tls.truststore.location	The location on disk of the trust store, in .jks format, used to confirm the authenticity of TLS/SSL servers that HBase Client RPC might connect to. This trust store must contain the certificate(s) used to sign the service(s) connected to. If this parameter is not provided, the default list of well-known certificate authorities is used instead.	CM_AUTO_TLS
fs.s3a.impl	The implementation to use for S3A FileSystem access.	org.apache.hadoop.hbase.oss.HBaseObjectStoreSemantics
hbase.rpc.tls.truststore.password	The password for the HBase Client RPC TLS/SSL Trust Store File. This password is not required to access the trust store; this field can be left blank. This password provides optional integrity checking of the file. The contents of trust stores are certificates, and certificates are public information.	CM_AUTO_TLS

New behavior:

Parameter name	Description	Default value
hbase.bucketcache.minfactor	The minimum percentage of BucketCache usage that should be targeted by the mass eviction process.	0.95

Parameter name	Description	Default value
hbase.server.netty.tls.client.auth.mode	The client authentication mode for mTLS authentication when TLS RPC is in use. Three modes are NEED, WANT and NONE.	NONE
ip_version	Specifies the IP version the service must use for network communication. • IPv4 - Uses IPv4 exclusively. • IPv6 - Uses IPv6 exclusively. • Dual-stack (IPv4 &IPv6) - Supports both IPv4 and IPv6, enabling communication over both protocols.	IPV4
hbase.client.netty.tls.enabled	Encrypt communication between clients and HBase RPC client mode using Transport Layer Security (TLS) (formerly known as Secure Socket Layer (SSL)).	FALSE
hbase.bucketcache.acceptfactor	The percentage of BucketCache usage that must trigger the eviction of blocks.	0.98
hbase.rpc.tls.truststore.location	The location on disk of the trust store, in .jks format, used to confirm the authenticity of TLS/SSL servers that HBase Client RPC might connect to. This trust store must contain the certificate(s) used to sign the service(s) connected to. If this parameter is not provided, the default list of well-known certificate authorities is used instead.	None
fs.s3a.impl	The implementation to use for S3A FileSystem access.	org.apache.hadoop.fs.s3a.S3AFileSystem
hbase.rpc.tls.truststore.password	The password for the HBase Client RPC TLS/SSL Trust Store File. This password is not required to access the trust store; this field can be left blank. This password provides optional integrity checking of the file. The contents of trust stores are certificates, and certificates are public information.	None

Summary: Add security headers to Thrift/HTTP server

Previous behavior:

Previously, security headers were absent in the responses of Thrift or HTTP servers when utilizing HBase-exposed services.

New behavior:

To address this, additional security headers are introduced in the responses of the HBase Thrift server when HTTP or HTTPS transport is enabled.

Apache JIRA: HBASE-27118

Behavioral Changes in HDFS

Behavioral changes denote a marked change in behavior from the previously released version to this version of HDFS.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

HDFS Offline Image Viewer (oiv) returns non-zero error code on failure

Summary:

The HDFS Offline Image Viewer `oiv` command now exits with a non-zero status code and displays an error message upon failure instead of returning code 0.

Previous behavior:

Previously, when the `oiv` command failed due to system-level issues (such as full disks or network errors), it exited with a status code of 0 (success), even if the output file was truncated or incomplete.

New behavior:

Now, the `oiv` command detects operational failures and exits with a non-zero status code and a descriptive error message, enabling users and automation scripts to reliably detect execution failures.

Cloudera Runtime 7.3.2.100 CHF 1

Improved error handling for HDFS Offline Image Viewer (oiv)

Summary:

The HDFS Offline Image Viewer (oiv) command has been updated to correctly report failures when converting an fsimage to CSV.

Previous behavior:

When the `oiv` command failed due to system-side issues (for example, local filesystem quota exhaustion, disk full, or networking issues), it would exit with code 0 (success). This resulted in truncated or incomplete output files without any indication of failure to the user or automation scripts.

New behavior:

The `oiv` command now correctly identifies when the output file is incomplete or when a system-level error occurs during processing. In such cases, the command will exit with a non-zero error code and provide an error message, allowing users and scripts to detect the failure.

Cloudera Runtime 7.3.2

There are no behavioral changes in this release.

Behavioral Changes in Hive

Functional adjustments and behavioral updates for Hive are introduced in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

Avro tables no longer fetch schema files over HTTP or HTTPS by default

Hive restricted how the Avro SerDe resolves table schemas from the avro.schema.url table property, to prevent server-side request forgery.

Previous behavior:

When avro.schema.url was not a Hadoop file system URI, Hive fetched the schema from any URL, including http://, https://, and file:// locations, using the identity of the HiveServer2, Metastore,

or Tez or MapReduce worker process. No scheme or host validation was performed, and no authorization check was made on the schema location, so a user could point an external Avro table at an arbitrary internal or cloud metadata endpoint and have the server fetch it on the user's behalf.

New behavior:

Hive now validates the scheme of avro.schema.url before fetching it. Schema-less URIs and configured file system schemes, such as hdfs, s3a, and abfs, continue to be allowed, and are controlled by the new hive.avro.schema.url.allowed.schemes property. The file:// and ftp:// schemes are always blocked. HTTP and HTTPS schema URLs are blocked by default. To restore the previous behavior for HTTP and HTTPS, set hive.avro.schema.url.remote.http.enabled to true and add the permitted hosts to hive.avro.schema.url.http.allowed.hosts. Hive also now enforces DFS URI authorization on file system schema paths at both DDL time and query time, including for tables created before this fix.

Apache Jira: HIVE-29671

Cloudera Runtime 7.3.2.10000 SP1

Summary: The default values for hive.server2.async.exec.threads and hive.server2.async.exec.wait.queue.size configurations increased

Previous behavior: The default values for the hive.server2.async.exec.threads and hive.server2.async.exec.wait.queue.size configurations were set to 100.

New behavior: Starting in Cloudera Manager 7.13.2.10000 and Cloudera Runtime 7.3.2.10000, the default values for the hive.server2.async.exec.threads configuration and the hive.server2.async.exec.wait.queue.size configuration are now 250. You can modify these configuration values in Cloudera Manager.

1. Go to the Cloudera Manager Clusters Hive service Configuration page.
2. Search for the Hive Service Advanced Configuration Snippet (Safety Valve) for hive-site.xml, click ⊕ and add:
   - Name hive.server2.async.exec.threads

     Value <Number of threads in the async thread pool>
   - Name hive.server2.async.exec.wait.queue.size

     Value <Size of the wait queue for async thread pool in HiveServer2>
3. Click Save.
4. Restart HIVE.

Summary: Increasing Hive Metastore comment length

Previous behavior:

The Hive Metastore schema explicitly restricted table and column comments to a maximum of 256 characters. When you attempted to enter or migrate comments longer than 256 characters, the process failed or truncated the text.

New behavior:

The Hive Metastore schema now consistently supports table and column comments up to 4000 characters across all supported backend databases. Hive Metastore configuration mappings and internal API validation checks reflect this new maximum length. Database migration scripts are provided to upgrade existing schemas automatically. You can now run a comment command on a table or column with up to 4000 characters by using Hive or Impala, and the backend database safely persists the data without truncation or errors.

Note: If you modified the comment character set or length before this upgrade, you must consider the following conditions:

- The comment column is based on the latin1 character set.

- The length of the comment column is less than 4000 characters.

Cloudera recommends backing up and taking a dump of the backend database before you perform the upgrade to ensure you can restore the data.

CDPD-98927

Cloudera Runtime 7.3.2.100 CHF 1

Summary:

Performance-related issues for HMS

Previous behavior:

Several performance-related issues for HMS occurred, including stale entries or inconsistent metadata in the database and exceptions during the update operation for multiple table column statistics or for multiple partition-level column statistics.

New behavior:

To resolve these issues, a backend database exclusive row lock is now introduced. The statistics importer tasks are coordinated so that only one importer task that obtains the lock runs at a time. The other tasks to import or update the statistics for the same table or partition continue to be in the pending queue until the current importer task completes and then releases the lock. By default, a blocked or pending importer task retries the lock 100 times in the 3 secs ~ 33 secs duration until a failure occurs. You can now configure the number of retries and the interval between retries using the following parameters that you can set on the Cloudera Manager Clusters Hive service Configuration tab :

- The hive.metastore.s4u.nowait.max.retries parameter allows you to set the number of retries to acquire the row lock immediately without waiting. The default value is 100.
- The hive.metastore.s4u.nowait.retry.sleep.interval parameter allows you to set the interval between the retries to acquire a row lock immediately. The default value is 300ms.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Hive, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see Behavioral Changes.

Summary:

Handling invalid date formats in to_date function.

The behavior of the to_date function has changed between Cloudera Runtime versions 7.1.7 SP2 and 7.1.9 when handling invalid date formats.

Previous behavior:

Valid dates (e.g., YYYY-MM-DD): Returned correct results.

Invalid dates: Returned random, unexpected dates instead of NULL.

New behavior:

Following the changes introduced inHIVE-28483:

Valid dates: Continue to return correct results.

Invalid dates (e.g., DD-MM-YYYY): Now return NULL.

Summary:

Metastore secondary connection pool size is now configurable

Previous behavior: The Metastore's secondary connection pool had a fixed size of 2. This often led to connection limitations, especially under heavy workloads.

New behavior:: You can now configure the metastore's secondary connection pool size using the property datanucleus.connectionPool.secondary.maxPoolSize. This lets you adjust the pool beyond its default of 2, preventing connection limitations and improving performance.

Summary:

New configuration available to disable the Partition Management Task

Previous behavior:

It was not possible to disable the PartitionManagementTask, which put a heavy load on the Cloudera Manager Metastore when managing a large number of tables and partitions. Customers had a tedious workaround by providing a pattern via metastore.partition.management.database.pattern / metastore.partition.management.table.pattern.

New behavior:

You can now set the metastore.partition.management.task.frequency configuration to 0 to disable the partition management task cluster-wide. This helps to reduce the load on the Cloudera Manager Metastore.

Apache Jira: HIVE-25324

Summary:

New default restriction for Iceberg table data file locations

Previous behavior:

In earlier versions, the hive.iceberg.allow.datafiles.in.table.location.only property was set to false by default. This allowed Hive to access and read Iceberg data files even if they were located outside of the specific table directory.

New behavior:

The default value for hive.iceberg.allow.datafiles.in.table.location.only is now true. This security enhancement ensures that only data files located within the table directory are accessible. If you attempt to read an Iceberg table that contains data files outside of its directory, Hive now returns an error. If your existing workflows rely on data files stored in external locations, you can disable this restriction by using the Cloudera Manager to set the property to false.

Summary:

Lineage information computation enabled by default

Previous behavior:

Previously, lineage information was computed only if specific hardcoded post-execution hooks were configured or if the deprecated HIVE_LINEAGE_INFO property was set to true.

New behavior:

Lineage information is now collected by default for all queries. However, the system only records and passes lineage to hooks for the specific query types defined in the HIVE_LINEAGE_STA TEMENT_FILTER property. By default, this includes CREATE_TABLE, CREATE_TABLE_AS_ SELECT,

CREATE_VIEW, CREATE_MATERIALIZED_VIEW, and LOAD.

Apache Jira: HIVE-28768

Summary:

Increased Batch Sizes for COMPUTE STATS

Previous behavior:

The `COMPUTE STATS` query previously failed on tables containing more than 5000 columns. This issue was specific to wide tables and could not be resolved by dropping and rerunning the query.

New behavior:

To resolve this, we enable the batch retrieval or insertion of the object metadata by default. The default value of the hive.metastore.direct.sql.batch.size property is changed from 0 to 1000, and the default value of the metastore.rawstore.batch.size property is changed from -1 to 500. After this change, `COMPUTE STATS` queries now run successfully on tables with more than 5000 columns.

Behavioral Changes in Cloudera Data Explorer (Hue)

Functional adjustments and behavioral updates for Cloudera Data Explorer (Hue) are introduced in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

Summary: Removal of automatic VACUUM from manual cleanup API

Previous behavior:

Previously, the Query Processor manual cleanup API (/cleanup/{startTime}) automatically ran VACUUM after deleting records. This behavior caused compatibility issues with MySQL, unexpected resource consumption, and database locks during disk space recovery.

New behavior:

The manual cleanup API no longer runs VACUUM automatically after deleting records. Database administrators must run VACUUM manually when necessary.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Cloudera Data Explorer (Hue), and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see Behavioral Changes.

There are no behavioral changes introduced in this release.

Behavioral Changes in Impala

Functional adjustments and behavioral updates for Impala are introduced in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

Summary: Disabling Impala AI functions in FIPS mode

Previous behavior:

The Impala AI functions ai_generate_text and ai_generate_text_default rely on Java Cryptography Extension KeyStore (JCEKS) or plain text for credential storage. Previously, these functions could run in Federal Information Processing Standards (FIPS) mode.

New behavior:

The ai_generate_text and ai_generate_text_default functions are disabled when running in FIPS mode. If you attempt to use these functions while in FIPS mode, the system displays a clear error message.

Summary: Adding a catch-up mode to EventProcessor

Previous behavior:

Previously, when Hive Metastore (HMS) event processing experienced a large lag or during a catalogd High Availability (HA) failover, the EventProcessor component performed full table reloads for table-level events. This could cause slow recovery times when processing large event backlogs.

New behavior:

Added a configurable catch-up mode by using the hms_event_catchup_threshold_s property for HMS event processing. This mode invalidates tables instead of performing expensive reloads when the event lag exceeds a configurable threshold, which improves recovery from large event backlogs.

Apache Jira: IMPALA-14230

Cloudera Runtime 7.3.2.100

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Impala, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see Behavioral Changes.

Summary:

Query cancellation during frontend planning

Previous behavior:

Previously, you could not cancel queries while they were in the analysis or planning stages. If a query was waiting for metadata to be loaded from the Catalog Server or the Hive Metastore (HMS), it would continue to run until the frontend planning process was complete. You had to wait for the query to reach the execution stage before a cancellation request could take effect.

New behavior:

You can now cancel queries during the frontend planning and metadata operation stages. If a query is interrupted by a user or a timeout while in the planning stage, the system triggers an interruption to the specific thread. This is particularly effective when planning is waiting on external metadata services. The cancellation process blocks until the frontend reaches an interruption point and returns to the backend to finalize the query status.

Summary:

Default value for fe_service_threads increased to improve concurrency

Previous behavior:

The default value for the fe_service_threads setting was 64.

New behavior:

Starting with Cloudera on premises 7.3.2, the default value is 128.

Summary:

Cleanup subdirectories in truncate/insert overwrite if recursing listing is enabled

Previous behavior:

Impala did not consistently delete files located in subdirectories of external tables during `TRUNCATE` and `INSERT OVERWRITE` operations, even when recursive listing was enabled. This led to leftover data in subdirectories after these operations, resulting in data corruption.

New behavior:

After this change, directories are also deleted in addition to (non-hidden) data files, with the exception of hidden and ignored directories. Now, setting DELETE_STATS_IN_TRUNCATE=false is no longer supported by default when truncating non-transactional tables; attempting this will result in an exception. If the old behavior is absolutely required, you can set the --truncate_external_tables_with_hms flag to false, but be aware that this will also reintroduce the bug that was fixed by this change.

Apache Jira: IMPALA-14189, IMPALA-14224

Summary:

Parquet late materialization behavior has changed

Previous behavior:

Parquet late materialization feature was disabled by default. You would use the parquet_late_materialization_threshold query option to set the minimum number of consecutive filtered rows required to trigger late materialization. The default value was -1. The feature was not supported for collection columns.

New behavior:

Parquet late materialization feature is enabled by default for all types including collections. The parquet_late_materialization_threshold is now set to 1 if the query option is greater than or equal to 0 and there is a collection value that can be skipped. Otherwise, the value is the same as the query option, which defaults to 20.

Apache Jira: IMPALA-3841

Summary:

TCP Keepalive is now enabled by default for client connections

Previous behavior:

TCP keepalive was disabled by default for client connections. Idle connections dropped by load balancers remained active in Impala, consuming service threads (fe_service_threads).

New behavior:

TCP keepalive is now enabled by default for all client connections, enhancing stability and availability. Impala is configured to check idle connections aggressively, every 10 minutes.

Apache Jira: IMPALA-14031

Summary:

Support for load-based routing in impala-proxy

Previous behavior:

The impala-proxy used a random selection policy to choose a coordinator. This approach did not consider the current load on each coordinator, which led to an uneven distribution of connections and potential performance bottlenecks.

New behavior:

The impala-proxy now uses load-based routing to decide which coordinator should handle a new session request. The Impala proxy directs the new session to the coordinator with the minimum calculated load. You can customize how this load is calculated using the following parameters:

- IMPALA_PROXY_COORDINATOR_LOAD_CPU_WEIGHT: Determines the weight applied to the current percentage of CPU utilization when calculating the coordinator's load.

- IMPALA_PROXY_COORDINATOR_LOAD_MEMORY_WEIGHT: Determines the weight applied to the current percentage of memory utilization when calculating the coordinator's load.

By adjusting these weights, you can tune the Impala proxy to prioritize CPU or memory headroom when routing new sessions.

Summary:

New catalogd flag for HMS event sync defaults

Previous behavior:

Previously, disabling event processing globally required you to manually set the impala.disableHmsSync property for every individual database and table.

New behavior:

The new disable_hms_sync_by_default flag now defines the global default for event processing. If set to true, Impala skips event processing for all tables and databases unless the impala.disableHmsSync property is explicitly set to false at the table or database level. The priority for checking the sync status is the table property, followed by the database property, and finally the global default flag.

Apache Jira: IMPALA-14085

Behavioral Changes in Apache Iceberg

Behavioral changes denote a marked change in behavior from the previously released version to this version of Iceberg.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

There are no behavioral changes in this release.

Behavioral Changes in Iceberg REST Catalog

Functional adjustments and behavioral updates for Cloudera Iceberg REST Catalog are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Behavioral Changes in Kafka

Functional adjustments and behavioral updates for Kafka are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

Kafka Connect Secrets Storage key wrap algorithm is configurable

Previous behavior:

The algorithm that Kafka Connect Secrets Storage uses to wrap the encryption key was hardcoded to AESWrap. You could not change this value in Cloudera Manager.

New behavior:

You can configure the wrap algorithm using the Kafka Connect Secrets Storage Secret Key Wrap Algorithm (kafka.connect.secret.key.wrap.algorithm) property in Cloudera Manager. The default remains AESWrap. On FIPS-enabled clusters where the Bouncy Castle FIPS jars are configured, Cloudera recommends setting this property to AESKWP. Other FIPS-approved algorithms supported by the configured security provider may also be used; AESKWP is the algorithm validated by Cloudera. Changing the wrap algorithm does not re-encrypt secret values stored in the secrets topic. If you change the value from the default and need to run the connect-secret-storage-migration tool, pass kafka.connect.secret.key.wrap.algorithm in the source topic configuration file, similar to kafka.connect.secret.global.password and kafka.connect.secret.pbe.salt.

HTTP metrics SSL enabled by default when Encrypt All Ports is on

Previous behavior:

The default value of SSL Encryption For HTTP Metric Reporter (kafka.http.metrics.ssl.enabled) was false on Kafka Broker and KRaft Controller roles.

New behavior:

When you enable the ENCRYPT_ALL_PORTS feature flag, the effective default becomes true. If you previously set a custom value, your setting is preserved. For information about the feature flag, see *Transport Layer Security Support and Custom Protocol Configuration*.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Kafka, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see Behavioral Changes.

Component-level custom Java home configuration removed

Previous behavior:

You could configure a component-specific Java home for Kafka Connect.

New behavior:

The component-level custom Java home configuration options are removed. Kafka Connect now uses the host-level java_home configuration. If you previously set a component-specific Java home for this service, verify the host-level java_home setting after upgrading.

The Kafka service is configured to use KRaft by default

Previous behavior:

Newly deployed Kafka services were configured to use ZooKeeper as their metadata store. The Kafka Metadata Store Service (metadata.store) Kafka property was set to ZooKeeper.

New behavior:

Newly deployed Kafka services are configured to use KRaft as their metadata store. The Kafka Metadata Store Service (metadata.store) Kafka property is set to KRaft.

High Watermark no longer advances when ISR is below MinISR

Previous behavior:

The High Watermark (HWM) advanced regardless of whether the in-sync replica (ISR) count was below min.insync.replicas. When producers used acks=1 or acks=0, messages were written to the leader and became consumable once the HWM advanced, even if the ISR had dropped below the minimum threshold. The min.insync.replicas setting only affected acks=all produce requests, blocking writes when ISR was insufficient, but did not prevent HWM advancement or consumer reads for acks=0/acks=1 messages.

New behavior:

The HWM no longer advances when the ISR count falls below min.insync.replicas. As a result, consumers are blocked from reading new messages in this condition, even if producers with acks =1 or acks=0 are still writing to the leader. This ensures data is only consumable when it meets the cluster's minimum durability requirements. If you use min.insync.replicas=2 or higher, you may see reduced consumer throughput when the ISR count drops below the configured minimum.

For more information, see KIP-966: Eligible Leader Replicas and KAFKA-15583.

Kafka protocol version is set automatically during upgrades

Previous behavior:

The inter.broker.protocol.version property for ZooKeeper-based clusters and the metadata.version property for KRaft-based clusters were not set automatically before an upgrade. Manually configuring these properties to the current protocol and metadata version was required before an upgrade.

New behavior:

During a cluster upgrade, Cloudera Manager now automatically sets the inter.broker.protocol.version property for ZooKeeper-based clusters and the metadata.version property for KRaft-based clusters. Manual configuration is no longer required.



Note:

Manually clearing these properties after the upgrade is still necessary. However, both properties are now available for direct configuration in Cloudera Manager. Using advanced configuration snippets is no longer required.

Behavioral Changes in Apache Knox

Functional adjustments and behavioral updates for Apache Knox are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Knox, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see Behavioral Changes.

Role-level alias management for Knox Gateway and IDBroker

Summary:

The alias management configuration has been moved from service-level to role-level, with dedicated configurations for each role.

Previous behavior:

Previously, both Knox Gateway and IDBroker aliases were saved using the shared service-level save_alias_command_input_password configuration and a single Save Alias command.

New behavior:

Now, each role has its own dedicated role-level configuration: gateway_save_alias_command_input for the Knox Gateway role and idbroker_save_alias_command_input for the IDBroker role. Two role-specific commands are available: Save Alias - Knox Gateway and Save Alias - IDBroker. For more information, see Saving aliases.

Behavioral Changes in Kudu

Functional adjustments and behavioral updates for Kudu are introduced in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

There are no behavioral changes in this release.

Behavioral Changes in Livy

Behavioral changes denote a marked change in behavior from the previously released version to this version of Livy.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

For encryption-on-the-wire overview and TLS guidance, see .

OPSAPS-72288: Spark 3 dependency extension for Livy for Spark 3

Previous behavior:

Livy for Spark 3 duplicated some Spark 3 paths in its own Cloudera Manager service descriptor, including the spark_pyspark_executable_path parameter.

New behavior:

Livy for Spark 3 uses a Spark 3 dependency extension in Cloudera Manager. Shared variables are inherited from the Spark 3 service, and you no longer maintain a separate spark_pyspark_executable _path value on the Livy service.

Summary: Livy server default JVM trust store password for BCFKS cacerts

Previous behavior:

Livy for Spark 3 did not set javax.net.ssl.trustStorePassword when the JVM default trust store was BCFKS, which could cause the Livy server to fail on FIPS 140-3 clusters.

New behavior:

When livy.ssl.defaultTrustStorePassword is set, the Livy server sets javax.net.ssl.trustStorePass word if it is not already defined. See .

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Livy, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see Livy Behavioral Changes.

There are no behavioral changes in this release.

Behavioral Changes in Oozie

Functional adjustments and behavioral updates for Oozie are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

Oozie capture-output will not fail for long input data

Previous behavior:

Previously, the Oozie capture-output failed if the input data exceeded the maximum allowed length.

New behavior:

Now, Oozie action works correctly even if the output data exceeds the maximum allowed length. Instead, Oozie reads only the first N characters, as defined by your configuration, from the output. To revert this behavior, set the oozie.servlet.CallbackServlet.fail.on.data.overflow and

oozie.action.fail.on.output.data.overflow properties to true in the Oozie Server Advanced Co nfiguration Snippet (Safety Valve) for oozie-site.xml advanced configuration snippet.

Behavioral Changes in Ozone

Behavioral changes for Cloudera Runtime 7.3.2. introduce functional adjustments and behavioral updates for Ozone.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

Introduction of DiskBalancer for Ozone DataNodes

Previous behavior:

Previously, Ozone DataNodes lacked an internal mechanism to balance data across multiple disks within a single node, leading to uneven disk utilization and storage inefficiency.

New behavior:

Each DataNode now includes a DiskBalancer to balance data evenly across its internal disks. This feature is disabled by default and can be enabled by setting hdds.datanode.disk.balancer.enabled configuration property to true. For more information, see Disk Balancer.

Clarification and separation of --all and --all-status options in ozone sh snapshot lsDiff command.

Previous behavior:

Previously, the `--all` option was renamed to `--all-status`, which unintentionally exposed an underlying `--all` option. This caused confusion as the distinction between removing pagination limits and filtering by job status was not clearly documented or implemented, leading to inconsistent behavior in scripts.

New behavior:

Now, the options are clearly distinguished and documented:

- `--all` (or `-a`): Removes the pagination limit and lists every matching result.
- `--all-status`: Lists all snapshot diff jobs regardless of their current status (example, queued, in_progress, done, failed, rejected).

Improved Ozone Recon bootup time by making failed delta task reinitialization asynchronous.

Previous behavior:

Previously, during startup, Ozone Recon would perform a blocking call to reinitialize any delta tasks that had failed in a previous run. This synchronous process delayed the overall bootup time of Recon, especially if there were multiple failed tasks to process.

New behavior:

Now, the reinitialization of failed delta tasks at startup is handled as an asynchronous event. By moving these calls to a background queue, Recon can complete its bootup process more quickly without waiting for the reinitialization tasks to finish.

Reserved space added to DataNode capacity tooltip

Previous behavior:

Previously, the capacity tooltip in the DataNodes tab did not display the reserved space, making it difficult to distinguish overall filesystem capacity from the actual capacity available to Ozone.

New behavior:

Now, the DataNode capacity tooltip now explicitly displays the Reserved Space value, clearly showing the difference between total filesystem capacity and storage allocated to Ozone.

Dedicated logging for Ozone Manager deletion services and SnapDiff workers

Previous behavior:

Previously, deletion service and Snapshot Diff (SnapDiff) worker logs were routed to the main Ozone Manager log file (ozone-om.log), making it difficult to isolate and troubleshoot garbage collection or snapshot activities.

New behavior:

Deletion-related logs are now separated from ozone-om.log. All deletion service logs are consolidated into a dedicated rolling log file (om-gc.log).

Ozone snapshot rename feature disabled by default

Previous behavior:

Previously, the snapshot rename feature was enabled by default, allowing users to rename snapshots despite potential instability that could cause snapshots to become unfindable.

New behavior:

Now, the snapshot rename feature is disabled by default to prevent data integrity issues. A new configuration property, ozone.om.snapshot.rename.allowed, controls this feature and defaults to false. To enable snapshot renames, administrators must explicitly set this property to true in the Ozone Manager configuration.

Ozone snapshot metadata now records successful defragmentation timestamps

Previous behavior:

Previously, the lastDefragTime field in snapshot metadata YAML files remained unpopulated, leaving administrators without a direct way to verify when a snapshot was last defragmented.

New behavior:

The lastDefragTime field is now automatically updated whenever a snapshot database defragmentation process completes successfully, providing clear visibility into snapshot maintenance status.

Locking and new checkpoint endpoint added to Ozone Manager bootstrap process

Previous behavior:

Previously, the Ozone Manager (OM) bootstrap process did not lock snapshotted RocksDB databases. Active transactions could occur on the snapshot database during bootstrap, creating a race condition that resulted in database corruption on follower nodes and bootstrap failures.

New behavior:

Now, the OM bootstrap process uses locks to prevent race conditions and protect snapshot database integrity. Additionally, clients running Apache Ozone version 2.2.0 or later now use a new checkpoint endpoint (/v2/dbCheckpoint), while older clients continue using the legacy endpoint (/dbCheckpoint). This functionality is enabled by default and can be controlled using the ozone.om.db.checkpoint.use.inode.based.transfer configuration property.

Ozone snapshot deep cleaning enabled by default

Previous behavior:

Previously, the ozone.snapshot.deep.cleaning.enabled configuration property defaulted to false. When a snapshot was deleted, keys marked for deletion were perpetually moved to the next active snapshot in the chain rather than being purged, preventing storage reclamation and causing unreferenced data to accumulate.

New behavior:

The ozone.snapshot.deep.cleaning.enabled property now defaults to true. Deleted keys progress through the snapshot chain until they are fully purged, allowing the system to reclaim storage automatically without manual configuration.

Explicit error messages for invalid Ozone bucket and volume names

Previous behavior:

Previously, creating an Ozone bucket or volume with an invalid name (such as one containing an underscore) generated a generic error message that did not specify why validation failed.

Example,

```
INVALID_BUCKET_NAME Invalid bucket name: bucket_1
```

New behavior:

The Ozone CLI now returns a detailed error message that explicitly identifies the unsupported character causing the validation failure.

Example,

```
INVALID_BUCKET_NAME bucket name has an unsupported character : _
```

Ozone create, write, and delete operations now require Read permission

Previous behavior:

Creating, writing, or deleting an Ozone key did not require READ permission. If a user lacked access, the system returned a WRITE or DELETE permission denied error.

New behavior:

The Ozone Manager now checks permissions using the getFileStatus() API before creating, writing, or deleting. Users must have READ permission (or SELECT permission for Hive resources when Ranger RMS is enabled) for these actions to succeed. Unauthorized requests now return a READ permission denied error. This aligns with AWS S3 behaviour, preventing unauthorized users from checking whether a resource exists.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

Renamed Ozone Manager Compaction Service Configuration Property

Previous behavior:

Previously, the property used to enable or disable the compaction service was ozone.compaction.service.enabled.

New behavior:

Now, to follow Ozone Manager (OM) naming conventions, the property is now ozone.om.compaction.service.enabled.

Migration of the `ozone debug recover` command to the `admin om lease recover` Ozone admin CLI command

The `ozone debug recover` command, used for lease recovery, now resides under the `admin om lease recover` Ozone admin CLI command.

Example,

```
# ozone admin om lease recover --path=ofs://<OZONE.SERVICE.ID>/
hsyncvol/hsyncbuck/hsync/File_0.txt
```

```
Lease recovery SUCCEEDED on ofs://<OZONE.SERVICE.ID>/hsyncvol/hs
yncbuck/hsync/File_0.txt
```

Changed functionality of the `--all` option in the `ozone sh snapshot lsDiff` command

Previous behavior:

Previously, the `--all` option was used to list snapshot diff jobs of any status.

New behavior:

Now, the `--all` option is used for listing limit. When the `--all` option is set, the listing limit is effectively set to Integer.MAX_VALUE. For the original `--all` behavior to list snapshot diff jobs, use the `--all-status` option instead.

Changed listing and deleting behavior at volume level

Previous behavior:

Previously, recursive listing and deleting a volume was supported using the `ozone fs` command that worked only if FSO bucket existed in volume.

New behavior:

Now, recursive listing and deleting a volume is supported only by using the `ozone shell` command and not the `ozone fs` command. The `ozone shell` command works for both OBS and FSO buckets inside a volume.

Example

- Recursive delete command

  ```
  # ozone fs -rm -R -skipTrash ofs://ozone1713158544/vol/
  ```

  Sample output

  ```
  24/04/15 13:12:47 WARN ozone.BasicRootedOzoneFileSystem: Rec
  ursive volume delete using ofs is not supported
  rm: Recursive volume delete using ofs is not supported. Inst
  ead use 'ozone sh volume delete -r -skipTrash -id <OM_SERVIC
  E_ID> <Volume_URI>' command
  ```

- Recursive list command

  ```
  # ozone fs -ls -R ofs://ozone1713158544/vol/
  ```

  Sample output

  ```
  24/04/15 13:12:53 WARN ozone.BasicRootedOzoneFileSystem: Rec
  ursive root/volume list using ofs is not supported
  ls: Recursive list root/volume using ofs is not supported. In
  stead use 'ozone sh key list <Volume_URI>' command
  ```

Migration of the ozone debug chunkinfo API to the ozone debug replicas chunk-info API

Previous behavior:

Previously, the ozone debug chunkinfo API returned the following JSON response:

```
{
  "KeyLocations" : [ [ {
    "Locations" : {
      "files" : [ "<block_path>" ],
      "pipelineID" : "<pipeline_id>"
    },
    "Datanode-HostName" : "<host_name>",
```

```
    "Datanode-IP" : "<host_ip>",
    "Container-ID" : <id>,
    "Block-ID" : <blockid>
  },{},{}
]]}
```

New behavior:

Now, the ozone debug chunkinfo API is relocated to the ozone debug replicas chunk-info API and returns the following JSON response:

```
{
  "volumeName" : "<vol>",
  "bucketName" : "<buck>",
  "name" : "<key>",
  "keyLocations" : [ [ {
    "datanode" : {
      "hostname" : "<host_name>",
      "ip" : "<host_ip>",
      "uuid" : "<uuid>"
    },
    "file" : "<block_path>",
    "blockData" : {
      "blockID" : {
        "containerID" : <id>,
        "localID" : <localid>,
        "blockCommitSequenceId" : <bcsid>
      },
      "size" : <size>,
      "chunks" : [ {
        "offset" : <offset>,
        "len" : <len>,
        "checksums" : [ "<checksum>" ],
        "checksumType" : "<ctype>",
        "bytesPerChecksum" : <bpc>
      } ]
    }
  }, {}, {}
]]}
```

As part of this API update, the following keys are updated:

Table 13: JSON key mapping for the relocated ozone debug chunkinfo API

Key name	Key in the previous JSON	Key in the current JSON
Block path	data["KeyLocations"][0][0]["Locations"]["files"][0]	data["keyLocations"][0][0]["file"]
Container ID	"Container-ID"	blockData.blockID.containerID
Block ID	"Block-ID"	blockData.blockID.localID
Hostname	"Datanode-HostName"	datanode.hostname

The S3 Gateway web UI is moved to a new port

Previous behavior:

The S3 Gateway web UI was accessible on ports 9878 for HTTP and 9879 for HTTPS, just like S3 buckets.

New behavior:

The S3 Gateway web UI resides on new ports to avoid conflicts with S3 buckets. The new ports are 19878 for HTTP and 19879 for HTTPS.

Behavioral Changes in Parquet

Functional adjustments and behavioral updates for Parquet are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

There are no behavioral changes in this release.

Behavioral Changes in Phoenix

Functional adjustments and behavioral updates for Phoenix are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

Summary: Phoenix is updated to comply with FIPS 140-3 cryptographic requirements

Previous behavior:

Phoenix used cryptographic algorithms that are not approved under FIPS 140-2 or FIPS 140-3. In FIPS-enabled cluster deployments, this caused the Phoenix Query Server to fail to start, JDBC connection failures due to non-FIPS-approved TLS cipher suites, and query execution errors.

New behavior:

Phoenix is updated to replace all non-FIPS-compliant cryptographic operations with FIPS-approved equivalents. Following are the key changes:

- Phoenix JDBC connections now accept only FIPS 140-3-approved TLS cipher suites. You must configure Phoenix clients to use FIPS-compliant cipher suites before connecting to a FIPS-enabled cluster.
- For clusters upgraded to Cloudera Runtime 7.3.2.10000 with FIPS mode enabled, you must re-encrypt any existing Phoenix system table credentials using the updated key derivation scheme before running Phoenix queries.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

There are no behavioral changes in this release.

Behavioral Changes in Apache Ranger

Behavioral changes denote a marked change in behavior from the previously released version to this version of Apache Ranger.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

READ permission now required for Ozone key create/delete operations

Summary:

READ permission is now required for Ozone key create/delete operations

Previous behavior:

Earlier, while performing file creation on Ozone, only create and write permissions were checked on the Ozone key resource. Similarly, for file deletion, only the delete permission was checked. No read permission was required for either operation.

With Ranger RMS enabled, Ozone filesystem's create/write access was correctly translated to Hive create/update permissions, and Ozone delete access was correctly translated to Hive Drop permissions; without needing any read/select grant.

New behavior:

Ozone now calls getFileStatus before executing create, write, or delete operations on keys. This getFileStatus call performs a READ permission check. If the user lacks read access to the Ozone key resource, the request is denied with a READ permission error, even if the user has the correct create, write, or delete permissions.

When Ranger RMS is enabled, users now additionally need select permission (the Hive equivalent of read) in their Hive policy to perform create or delete operations on the underlying Ozone filesystem.

Ranger denies ALTER TABLE operations on tables with active row filter or column masking policies

Summary:

Ranger denies ALTER TABLE operations on tables with active row filter or column masking policies.

Previous behavior:

Earlier, when a row filter or column masking policy was active on a Hive table, the ALTER TABLE rename operation was permitted. As a result, the existing row filter and column masking policies become ineffective for the renamed table.

New behavior:

Currently, the ALTER TABLE operation is now denied when the source table has one or more active row filter or column masking policies. This change ensures that security policies remain valid and consistently applied to the intended table, regardless of any rename attempts.

GraalJS engine hardening for _expression policy conditions

Summary:

The GraalJS engine used to evaluate _expression policy conditions in Apache Ranger has been hardened for security.

Previous behavior:

Earlier, the GraalJS engine was configured with:

- polyglot.js.allowHostAccess=true — allowed unrestricted access to Java host objects from scripts
- polyglot.js.nashorn-compat=true — enabled Nashorn-style Java interop globals (Java.type(), java.*, Packages.*)

This permitted _expression policy conditions to use constructs such as:

```
Java.type('java.lang.Runtime')
java.lang.Runtime.getRuntime()
Packages.java.lang.String
new java.io.File('/tmp/test')
ctx.getClass().getClassLoader()
```

New behavior:

Currently, both allowHostAccess and nashorn-compat settings have been removed. The JavaScript constructs listed above will no longer work and will silently return null. When a blocked script is evaluated, evaluateScript() catches the exception and returns null. Depending on the policy structure, this could silently flip an ALLOW decision to DENY or vice versa.

If you are using Java class references directly in _expression conditions, you must rewrite them using the documented Ranger helper API. If a required capability is missing from the helper API, a JIRA should be filed to request a new helper method.

Enforcement of unauthenticated access properties

Summary:

The ranger.admin.allow.unauthenticated.access and ranger.admin.allow.unauthenticated.download. access properties are now enforced regardless of whether Kerberos authentication is configured on Ranger Admin. The default values remain unchanged.

Previous behavior:

Earlier, these settings were applied only when Ranger Admin was configured with Kerberos authentication. In non-Kerberos deployments, plugins and clients could access non-secure endpoints (to download policies, tags, roles, userstore, or to grant/revoke) without these properties being evaluated.

New behavior:

Currently, both properties are now honored consistently in all deployment modes (Kerberos and non-Kerberos). This is a breaking change for deployments where all of the following are true:

1. Ranger Admin is not configured with Kerberos authentication.
2. Ranger Admin is not configured to allow unauthenticated access.
3. Plugins/clients use the non-secure endpoints to download policies/tags/roles/userstore or to grant/revoke.

In affected deployments, take one of the following actions:

- Configure plugins/clients to use the secure endpoints. For example, set ranger.plugin.<serviceType>.forceSecureEndpointAccess=true.
- Configure Ranger Admin to allow unauthenticated access by setting ranger.admin.allow.unauthenticated.access=true (and/or ranger.admin.allow.unauthenticated.download.access=true).

Updated permission error response for unauthorized Ozone key deletion

Previous behavior:

Previously, attempting to delete an Ozone key without sufficient privileges returned a DELETE permission denied error.

New behavior:

The Ozone Manager now checks permissions through the getFileStatus() API during key deletion. Consequently, unauthorized delete requests return a READ permission denied error instead of a DELETE error. This aligns with AWS S3 behavior to prevent unauthorized users from verifying whether a resource exists.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.0

The behavioral changes for Apache Ranger in Cloudera Runtime 7.3.2.0 include all cumulative updates from previous releases (such as 7.3.1.x). This version specifically incorporates changes introduced in Cloudera Runtime 7.3.1.100 through 7.3.1.706 alongside the following functional adjustments. For a complete list, see Behavioral Changes.

Summary:

The new column authorization optimization property in the Ranger-HBase plugin changes Ranger audit behavior, when enabled. There is no behavioral change if the property is disabled.

Note: There may not be any performance increase if the user has explicit denials at the column level for any of the column families or the user is not fully authorized for the column families.

Previous behavior:

Earlier, there were more audit entries when the service configuration for optimization was not present and, hence, not enabled.

For example, the audit behavior for the `hbase ltt -tn multitest -families cf0,cf1 -write 3:2:1 -multiput -num_keys 1 -num_regions_per_server 1` command was as follows:

15 audit entries

New behavior:

Currently, there are fewer audit entries when the service configuration for optimization is enabled.

For example, the audit behavior for the `hbase ltt -tn multitest -families cf0,cf1 -write 3:2:1 -multiput -num_keys 1 -num_regions_per_server 1` command is as follows:

2 audit entries (multitest/cf1 and multitest/cf2)

Summary:

The following service configurations have been added to a new place in the Ranger Admin Web UI:

- Policy Download Users (policy.download.auth.users)
- Tag Download Users (tag.download.auth.users)
- Service Admin Users (service.admin.users)
- Service Admin Groups (service.admin.groups)
- Superusers (ranger.plugin.super.users)
- Superuser Groups (ranger.plugin.super.groups)
- Userstore Download Users (userstore.download.auth.users)

Previous behavior:

Previously, the above service configurations were available under the Add New Custom Configurations section.

New behavior:

Now, the above service configurations are directly available under the Config Properties section. Additionally, the configurations are added as dropdowns, where you can select the users and groups.

Also, after you upgrade to Cloudera Runtime 7.3.2.0 from any previous release, your existing configurations will be shifted to the Config Properties section with values.

Summary:

Policy resources in the Ranger Admin UI are being added using React JS instead of Backbone JS

Previous behavior:

Earlier with Backbone JS, when you copied and pasted resource values containing commas or spaces (for example, in Hive policy resources: database1, database2), the UI automatically split them into separate values — database1 and database2. The same behaviour applied to space-separated values. Because of this, you were not allowed to enter resource names containing commas and spaces, and this limitation affected all service policy resources.

New behavior:

After upgrading from Backbone JS to React JS, this restriction has been removed. Now, React JS treats pasted values with commas or spaces as a single entry. Hence, you can no longer paste multiple values at once; you must manually add each resource value.

Summary:

Hive authorization from Ranger for Alter Table Rename command does not require CREATE database permission on the database where the renamed table will be created.

Previous behavior:

Earlier, whenever Alter Table Rename command was used across databases in Hive, authorization from Ranger required CREATE database permission for the user on the target database in which the renamed table was created.

New behavior:

Now, whenever Alter Table Rename command is used across databases in Hive, authorization from Ranger does not check for CREATE database permission for the user on the target database in which the renamed table will be created.

Summary:

Added support for multiple columns policy creation in Ranger for Grant/Revoke request.

Previous behavior:

Previously, when a request with multiple columns, such as GRANT SELECT (col1, col2, col3, col4, col5, col6, col7, col8, col9, col10) ON TABLE demo.data5 TO ROLE testrole_09289898, is executed in Impala, it results in the creation of a separate grant policy for each column in Ranger.

New behavior:

Now, a request with multiple columns results in creation of a single policy for Grant request for all the columns in Ranger. Same is true for Revoke request.

Behavioral Changes in Ranger KMS

Behavioral changes denote a marked change in behavior from the previously released version to this version of Ranger KMS.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.0

Behavioral Changes in Schema Registry

Functional adjustments and behavioral updates for Schema Registry are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

Legacy TLS protocol and cipher settings replaced when TLS advanced control is enabled

Previous behavior:

You configured TLS protocol versions and cipher suites for the Schema Registry REST port using schema.registry.ssl.supportedProtocols, schema.registry.ssl.excludedProtocols, schema.registry.ssl.supportedCipherSuites, and schema.registry.ssl.excludedCipherSuites properties in Cloudera Manager.

New behavior:

When you enable the TLS_ADVANCED_CONTROL feature flag, Cloudera Manager hides the legacy supported and excluded properties and replaces them with schema.registry.ssl.allowedProtocols and schema.registry.ssl.allowedCipherSuites. These properties default to the centralized Cloudera Manager TLS settings. Custom exclude lists no longer apply when the flag is enabled. Cloudera Manager does not automatically migrate custom values. Review your TLS configuration before you enable the flag.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

Component-level custom Java home configuration removed

Previous behavior:

You could configure a component-specific Java home for Schema Registry.

New behavior:

The component-level custom Java home configuration options are removed. Schema Registry now uses the host-level java_home configuration. If you previously set a component-specific Java home for this service, verify the host-level java_home setting after upgrading.

Schema Registry now defaults to IPv4-only communication

The default value of the schema.registry.additional.java.options configuration parameter was updated to set the IP protocol to IPv4.

If you changed the default value of this parameter before upgrading, the new default value is not applied on upgrade. You can apply it manually after the upgrade.

Behavioral Changes in Solr

Functional adjustments and behavioral updates for Solr are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

Summary: FIPS 140-3 support for Solr

Previous behavior:

Solr did not support FIPS 140-3 compliance.

New behavior:

Solr supports FIPS 140-3 compliance on FIPS-enabled Cloudera Base on premises clusters. You must configure the required JVM flags and environment variables for the Solr service and all indexing clients, including the MapReduce Indexer Tool, CrunchIndexer Tool, HBase Indexer, and jobs submitted by using spark-submit, to ensure correct operation.

Summary: Truststore password configuration for MapReduce and CrunchIndexer jobs on FIPS-enabled clusters

Previous behavior:

You passed the truststore password to MapReduce Indexer Tool and CrunchIndexer Tool jobs by using the -Djavax.net.ssl.trustStorePassword JVM system property in job submission scripts.

New behavior:

You must set the truststore password by configuring the default_truststore_pwd parameter in the YARN service configuration in Cloudera Manager. Passing the password by using the -Djavax.net.ssl.trustStorePassword flag is no longer supported for these jobs on FIPS-enabled clusters.

Summary: ZooKeeper client configuration for Solr CLI on FIPS-enabled clusters

Previous behavior:

No specific ZooKeeper client environment variables were required for solrctl operations on FIPS-enabled clusters.

New behavior:

You must set the ZKCLI_JVM_FLAGS environment variable with the required JAAS configuration and ZooKeeper ACL provider before you run `solrctl` commands on a FIPS-enabled cluster. Do not set the *SOLR_FIPS_ENABLED* or *SOLR_ZK_SSL_CONNECTION_ENABLED* environment variables for `solrctl` operations.

Summary: Remote streaming and stream body disabled by default

Previous behavior:

The Cloudera-shipped Solr configuration templates enabled remote streaming (enableRemoteStreaming="true") by default, allowing requests to fetch content from remote stream.url or stream.file locations.

New behavior:

Remote streaming (enableRemoteStreaming) and stream body (enableStreamBody) are now disabled by default for newly created Solr collections. Existing collections are not modified, including during an upgrade. To disable these features for an existing Solr collection, complete the following steps:

1. For a given Solr collection, locate and update the corresponding solrconfig.xml file by removing the enableRemoteStreaming="true" flag completely from the file.

   Note: If the collection has a configoverlay.json file that defines the enableRemoteStreaming flag, it must also be updated before uploading the configuration to ZooKeeper.
2. Upload the updated configuration to ZooKeeper using the `solrctl` command. For example:

   ```
   solrctl config --upload admin_collection1_instancedir /root/
   admin_collection1_configs/ -overwrite
   ```
3. Reload the updated Solr collection. For example:

   ```
   solrctl collection --reload admin_collection1
   ```
4. To verify that the configuration has been successfully updated for a given Solr collection, download the corresponding configuration from ZooKeeper and check the enableRemoteStreaming flag in solrconfig.xml. For example:

   ```
   solrctl instancedir --get admin_collection1_instancedir /tmp/1
   ```

Behavioral Changes in Spark

Behavioral changes denote a marked change in behavior from the previously released version to this version of Spark.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

CDPD-98799: Rebase on Apache Spark 3.5.8

Previous behavior:

Bundled Spark in Cloudera Runtime 7.3.2.0 was based on Apache Spark 3.5.4.

New behavior:

From 7.3.2.10000 SP1, bundled Spark is based on Apache Spark 3.5.8. Even maintenance rebases can affect application code or defaults; review Upgrading Apache Spark 3.5.4 to Spark 3.5.8, Migrating Spark applications, and the Apache Spark 3.5.8 documentation before production rollout.

When you enable Encrypt all ports in Cloudera Manager, Cloudera Manager applies the Spark property changes in this section automatically. For an overview, see .

OPSAPS-76333: Spark authentication and network encryption when Encrypt all ports is enabled

Previous behavior:

spark.authenticate and spark.network.crypto.enabled were not enabled automatically when you turned on cluster-wide port encryption.

New behavior:

When Encrypt all ports is active, Cloudera Manager sets spark.authenticate and spark.network.crypto.enabled to secure Spark RPC with AES encryption. See .

OPSAPS-75940: Spark IO encryption key size and algorithm defaults in Cloudera Manager

Previous behavior:

Cloudera Manager exposed spark.io.encryption.keySizeBits with a default of 128 and did not surface spark.io.encryption.keygen.algorithm on the Spark configuration page.

New behavior:

Cloudera Manager sets the default spark.io.encryption.keySizeBits to 256 and the default spark.io .encryption.keygen.algorithm to HmacSHA256. These defaults apply when you enable optional on-disk IO encryption.

CDPD-94374, OPSAPS-75458: HTTPS-only Spark History Server and Spark UI when Encrypt all ports is enabled

Previous behavior:

When TLS was enabled, Jetty could still bind an HTTP port for the History Server or Spark UI unless you configured HTTPS-only listeners manually.

New behavior:

When Encrypt all ports is active and TLS is enabled, Cloudera Manager sets spark.ssl.historyServer. disableHttpPort and spark.ssl.ui.disableHttpPort to true. Jetty starts with HTTPS only. If TLS is not enabled, these properties are ignored.

CDPD-87464, OPSAPS-76633: spark.cloudera.auto.disable.insecure.spark.ui.enabled

Previous behavior:

The Spark UI could start without SSL even when insecure UI access was not intended.

New behavior:

When this property is enabled, Spark disables the UI if SSL is not active. The parameter is off by default and is enabled automatically when Encrypt all ports is active.

CDPD-87464, OPSAPS-76633: YARN auto-certificate for Spark UI in cluster mode

Previous behavior:

Spark did not use YARN ResourceManager ApplicationMaster certificates for the Spark UI in cluster mode unless you configured certificate policy manually.

New behavior:

When spark.cloudera.ui.yarn.autoCertificate.clusterMode.enabled is enabled, Spark can use a certificate issued to the ApplicationMaster for the YARN RM Proxy. This requires yarn.resourc emanager.application-https.policy set to LENIENT or STRICT in yarn-site.xml. When Encrypt all ports is active, Cloudera Manager sets the YARN policy to LENIENT automatically.

This feature applies only in cluster deploy mode. It does not apply in client mode because the driver runs on the submission host and the UI starts before YARN provides a certificate.

Summary: Default JVM trust store password for BCFKS cacerts

Previous behavior:

Spark did not set javax.net.ssl.trustStorePassword when the JVM default trust store was BCFKS, which could cause Spark History Server or jobs to fail on FIPS 140-3 clusters.

New behavior:

When spark.ssl.defaultTrustStorePassword is set, Spark supplements driver and executor JVM options and applies the password when Spark History Server creates a SecurityManager. See .

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Spark, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see Spark Behavioral Changes.

Summary: Enhanced the security of Spark by implementing new default configuration settings

Previous behavior:

1. spark.ui.enabled=true: possible initiation of an HTTP service that can be accessed from external hosts.
2. spark.io.encryption.keySizeBits=128: the default Spark keySizeBits

New behavior:

1. spark.ui.enabled=false: preventing initiation of an HTTP service that can be accessed from external hosts.
2. spark.io.encryption.keySizeBits=256: the default Spark keySizeBits has been increased from 128 to 256

Behavioral Changes in Spark Atlas Connector

Behavioral changes denote a marked change in behavior from the previously released version to this version of Spark Atlas Connector.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

Cloudera Runtime 7.3.2 introduces functional adjustments, behavioral updates for Spark Atlas Connector, and includes all service packs and cumulative hotfixes from 7.3.1.100 through 7.3.1.706. For a comprehensive record of all functional adjustments in Cloudera Runtime 7.3.1.x, see Spark Atlas Connector Behavioral Changes.

There are no behavioral changes in this release.

Behavioral Changes in Sqoop

Functional adjustments and behavioral updates for Sqoop are introduced in Cloudera Runtime 7.3.2, its service packs and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

There are no behavioral changes in this release.

Behavioral Changes in Streams Messaging Manager

Functional adjustments and behavioral updates for Streams Messaging Manager are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

Legacy TLS protocol and cipher settings replaced when TLS advanced control is enabled

Previous behavior:

The REST Admin Server and UI Server roles used four supported and excluded TLS protocol and cipher properties in Cloudera Manager (streams.messaging.manager.ssl.supportedProtocols, streams.messaging.manager.ssl.excludedProtocols, streams.messaging.manager.ssl.supportedCipherSuites, streams.messaging.manager.ssl.excludedCipherSuites, and the corresponding streams.messaging.manager.ui.ssl.* equivalents).

New behavior:

When you enable the TLS_ADVANCED_CONTROL feature flag, Cloudera Manager hides the legacy supported and excluded properties on both the REST Admin Server and the UI Server, and replaces them with streams.messaging.manager.ssl.allowedProtocols, streams.messaging.manager.ssl.allowedCipherSuites, streams.messaging.manager.ui.ssl.allowedProtocols, and streams.messaging.manager.ui.ssl.allowedCipherSuites. These properties default to the centralized Cloudera Manager TLS settings. Custom exclude lists no longer apply when the flag is enabled. Cloudera Manager does not automatically migrate custom values.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

Component-level custom Java home configuration removed

Previous behavior:

You could configure a component-specific Java home for Streams Messaging Manager.

New behavior:

The component-level custom Java home configuration options are removed. Streams Messaging Manager now uses the host-level java_home configuration. If you previously set a component-specific Java home for this service, verify the host-level java_home setting after upgrading.

Default JMX settings changed to restrict connections to localhost

Previous behavior:

The default value of the SMM_JMX_OPTS Cloudera Manager configuration option was -Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.authenticate=false -Dcom.sun.management.jmxremote.ssl=false, which allowed unrestricted, unauthenticated JMX connections.

New behavior:

The default value of SMM_JMX_OPTS has been changed to restrict JMX connections to localhost only and enable SSL. If you previously customized SMM_JMX_OPTS, your custom value is preserved on upgrade. To revert to open JMX, update SMM_JMX_OPTS in Cloudera Manager.

Streams Messaging Manager now defaults to IPv4-only communication

A new argument was added to SMM_JVM_PERF_OPTS that sets the IP protocol to IPv4 by default.

If you changed the default value of this parameter before upgrading, the new default value is not applied on upgrade. You can apply it manually after the upgrade.

Streams Messaging Manager UI Migration to Java

The Streams Messaging Manager UI service is migrated from a NodeJS runtime to a Java-based server. This change addresses security vulnerabilities associated with NodeJS dependencies and aligns Streams Messaging Manager with the centralized dependency management of the platform.

As a result of this migration, the following changes apply:

- Runtime environment

  The Streams Messaging Manager UI service now runs on the JVM. Configuration for the runtime environment is now managed via the SMM_JAVA_OPTS environment variable.
- TLS configuration

  TLS configuration moved from OpenSSL-style parameters to standard Java JSSE configuration. New Cloudera Manager parameters manage TLS protocols and cipher suites:

  - streams.messaging.manager.ui.ssl.supportedCipherSuites
  - streams.messaging.manager.ui.ssl.excludedCipherSuites
  - streams.messaging.manager.ui.ssl.supportedProtocols
  - streams.messaging.manager.ui.ssl.excludedProtocols
- Configuration migration

  During upgrade, Cloudera Manager attempts to automatically migrate existing TLS settings (including those found in the NODE_OPTIONS environment variable within safety valves) to the new Java-based configuration parameters. However, manual verification is strongly recommended.
- Safety valves

  Any properties previously set in the Streams Messaging Manager UI Server Environment Advanced configuration Snippet (Safety Valve) using NODE_OPTIONS that are not related to TLS must be manually translated to their Java equivalents (if applicable) and set using SMM_JAVA_OPTS.

Behavioral Changes in Streams Replication Manager

Functional adjustments and behavioral updates for Streams Replication Manager are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

Legacy TLS protocol and cipher settings replaced when TLS advanced control is enabled

Previous behavior:

You configured TLS protocol versions and cipher suites for the Streams Replication Manager Service role using streams.replication.manager.ssl.supportedProtocols, streams.replication.manager.ssl.excludedProtocols, streams.replication.manager.ssl.supportedCipherSuites, and streams.replication.manager.ssl.excludedCipherSuites properties in Cloudera Manager.

New behavior:

When you enable the TLS_ADVANCED_CONTROL feature flag, Cloudera Manager hides the legacy supported and excluded properties for the Streams Replication Manager Service role and replaces them with streams.replication.manager.ssl.allowedProtocols and streams.replication.manager.ssl.allowedCipherSuites. These properties default to the centralized Cloudera Manager TLS settings. Custom exclude lists no longer apply when the flag is enabled. Cloudera Manager does not automatically migrate custom values.

Metrics reporter SSL enabled by default when Encrypt All Ports is on

Previous behavior:

The default value of SSL Encryption for the Metrics Reporter (metrics.jetty.server.ssl.enabled) on the Streams Replication Manager Driver role and SSL Encryption For HTTP Metric Reporter (streams.replication.manager.service.metric.ssl.enabled) on the Streams Replication Manager Service role was false.

New behavior:

When you enable the ENCRYPT_ALL_PORTS feature flag, the effective default for both properties becomes true if you have not customized them. If you previously set custom values, your settings are preserved.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

Component-level custom Java home configuration removed

Previous behavior:

You could configure a component-specific Java home for Streams Replication Manager.

New behavior:

The component-level custom Java home configuration options are removed. Streams Replication Manager now uses the host-level java_home configuration. If you previously set a component-specific Java home for this service, verify the host-level java_home setting after upgrading.

Default JMX settings changed to restrict connections to localhost

Previous behavior:

The default value of the SRM_JMX_OPTS Cloudera Manager configuration option was -Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.authenticate=false -Dcom.sun.management.jmxremote.ssl=false, which allowed unrestricted, unauthenticated JMX connections.

New behavior:

The default value of SRM_JMX_OPTS has been changed to restrict JMX connections to localhost only and enable SSL. If you previously customized SRM_JMX_OPTS, your custom value is preserved on upgrade. To revert to open JMX, update SRM_JMX_OPTS in Cloudera Manager.

Public Kafka Connect endpoints removed from SRM REST server

The public Kafka Connect endpoints are removed from the SRM REST server. Previously, these endpoints allowed users to interact with the internal Kafka Connect cluster inside the SRM instance, including starting arbitrary connectors and modifying internal connectors. These endpoints were undocumented and not part of the official SRM API. They are no longer available.

Streams Replication Manager now defaults to IPv4-only communication

A new argument was added to SRM_JVM_PERF_OPTS that sets the IP protocol to IPv4 by default.

If you changed the default value of this parameter before upgrading, the new default value is not applied on upgrade. You can apply it manually after the upgrade.

Change in internal topic filtering logic

Summary:

The logic that identifies and filters internal topics in Streams Replication Manager has changed. This enables the replication of topics that appear to be internal but are not truly internal to Kafka and Streams Replication Manager, reducing the risk of unintentionally excluding user topics from replication.

Previous behavior:

Topics were filtered from replication by the `DefaultReplicationPolicy` and `IdentityReplicationPolicy` policies if the topic name ended with [***SEPARATOR***]internal. For example, .internal. In addition, the default deny list regex pattern was the following:

```
.*[\\-\\.]internal, .*\\.replica, __.*
```

New behavior:

Topics are now filtered from replication by the `DefaultReplicationPolicy` and `IdentityReplicationPolicy` policies if:

- Name starts with mm2 and ends with [***SEPARATOR***]internal
- Name ends with [***SEPARATOR***]checkpoints[***SEPARATOR***]internal

In addition, the default deny list regex pattern is now the following:

```
mm2.*\\.internal, .*\\.replica, __.*
```

Internal Kafka topics that start with a dot (.) or two underscores (__) are continued to be filtered from replication by default.

Note: The `ReplicationPolicy` interface also changed. Previously, topics were filtered from replication if their names ended with .internal or -internal. From now on, names ending with -internal are no longer filtered. In addition, topics are filtered from replication if:

- Name starts with mm2 and ends with .internal
- Name ends with .checkpoints.internal

Internal Kafka topics that start with a dot (.) or two underscores (__) are continued to be filtered from replication by default. If you have a custom policy that implements the interface, you might need to update it to correctly filter internal topics.

Behavioral Changes in YARN

Functional adjustments and behavioral updates for YARN and YARN Queue Manager are introduced in Cloudera Runtime 7.3.2, its service packs, and cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

Summary: Truststore password handling changed for FIPS 140-3 environments

In FIPS 140-3 mode, the default JVM truststore uses BCFKS format, which requires a password to be provided – unlike JKS, which assumes a default password. To preserve the auto-TLS experience and prevent MapReduce task failures, YARN now automatically injects the DEFAULT_TRUS TSTORE_PWD environment variable into containers when a FIPS environment is detected. The value of this variable is passed as the javax.net.ssl.trustStorePassword JVM property at task startup.

No manual configuration of /etc/default/cloudera-scm-server or truststore password settings are required. The default password value (changeit) is used and thus behavior is configurable using Cloudera Manager.

Summary: UpscaleServiceV2 for YARN autoscaling

Previous behaviour – When numerous Application Masters (AMs) are active, the previous (v1) autoscaling mechanism encounters a key limitation. Because it evaluates pending requests strictly against total available cluster capacity, a queue requiring 25% additional resources will fail to trigger scaling if 50% free capacity exists in other parts of the cluster. Consequently, the affected queue becomes frozen due to unfulfilled resource requests.

New behaviour – A new component, UpscaleServiceV2, is now introduced. It does not rely only on available cluster resources and it accounts for pending applications and AM constraints. Because it focuses on pending resource requests without regard to existing capacity, queue configurations do not impact scaling determinations. UpscaleServiceV2 is capable of initiating scale-out operations even when unallocated resources remain available in the cluster. The following new configuration properties are now added and can be configured through the Cloudera Manager UI:

Table 14: New configuration properties

Property	Default	Description
yarn.cluster.scaling.upscale.version	1	Set to 2 to enable UpscaleServiceV2
yarn.cluster.scaling.upscale.am.multiplier	2.0	Buffer multiplier for AM capacity headroom
yarn.cluster.scaling.upscale.threshold.cpu	0	Minimum pending CPU before scaling triggers
yarn.cluster.scaling.upscale.threshold.memory	0	Minimum pending memory before scaling triggers
yarn.cluster.scaling.upscale.threshold.apps	0	Minimum pending applications before scaling triggers

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

There are no behavioral changes introduced in this release.

Behavioral Changes in ZooKeeper

Behavioral changes denote a marked change in behavior from the previously released version to this version of ZooKeeper.

Cloudera Runtime 7.3.2.20000 SP2

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.10000 SP1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2.100 CHF 1

There are no behavioral changes in this release.

Cloudera Runtime 7.3.2

ZooKeeper authentication enforcement

Previous behavior:

ZooKeeper client connections were accepted even when they did not authenticate. Anonymous clients could open sessions with the ZooKeeper ensemble, as long as they could reach the client port (for example, 2181). This meant that components and tools that did not use ZooKeeper authentication could still connect and operate, as long as ACLs and network controls (such as firewalls) allowed it.

New behavior:

When ZooKeeper authentication enforcement is enabled, the ZooKeeper server only accepts connections from clients that successfully authenticate (for example, using Kerberos/SASL). Components and clients that do not authenticate, or that fail authentication, can no longer establish a session with ZooKeeper and therefore cannot connect or operate through the unsecured client port. This change hardens ZooKeeper security, but requires that all ZooKeeper#using components in the deployment be updated or configured to use authentication before enforcement is turned on.

Summary: The default values for the following configuration items are updated

Previous behavior:

Parameter name	Description	Default value
ip_version	Specifies the IP version the service must use for network communication. • IPv4 - Uses IPv4 exclusively. • IPv6 - Uses IPv6 exclusively. • Dual-stack (IPv4 &IPv6) - Supports both IPv4 and IPv6, enabling communication over both protocols.	-
sessionRequireClientSASLAuth	ZooKeeper configuration to enforce Simple Authentication and Security Layer (SASL) authentication.	-

New behavior:

Parameter name	Description	Default value
ip_version	Specifies the IP version the service must use for network communication. • IPv4 - Uses IPv4 exclusively. • IPv6 - Uses IPv6 exclusively. • Dual-stack (IPv4 &IPv6) - Supports both IPv4 and IPv6, enabling communication over both protocols.	IPV4
sessionRequireClientSASLAuth	ZooKeeper configuration to enforce Simple Authentication and Security Layer (SASL) authentication.	false

Release Matrix

You must be familiar with the versions of all the service packs and cumulative hotfixes for Cloudera Runtime 7.3.2.

Runtime Release Build Number	Cumulative Hotfix/ Service Pack	GBN	Release Date	Download URL
7.3.2.20000-258	7.3.2.20000 SP2	83281140	September 26, 2026	https://archive.cloudera.com/p/cdh7/7.3.2.20000/
7.3.2.10000-317	7.3.2.10000 SP1	82216952	August 28, 2026	https://archive.cloudera.com/p/cdh7/7.3.2.10000/
7.3.2.100-296	7.3.2.100 CHF 1	81274879	August 06, 2026	https://archive.cloudera.com/p/cdh7/7.3.2.100/
7.3.2.0-957	7.3.2	77083870	March 31, 2026	https://archive.cloudera.com/p/cdh7/7.3.2/

Cloudera Runtime Component Versions

List of the official component versions for Cloudera Runtime. To know the component versions for compatibility with other applications, you must be familiar with the latest component versions in Cloudera Runtime. You must also be aware of the available Technical Preview components and use them only in a testing environment.

Cloudera Runtime 7.3.2.20000

List of the official component versions for Cloudera Runtime 7.3.2.20000.

Apache Components

Component	Version
Apache Arrow	15.0.0.7.3.2.20000-258
Apache Atlas	2.4.0.7.3.2.20000-258
Apache Calcite	1.33.0.7.3.2.20000-258
Apache Avatica	1.26.0.7.3.2.20000-258
Apache Avro	1.11.3.7.3.2.20000-258
Apache Hadoop (Includes HDFS, MapReduce and YARN)	3.4.2.7.3.2.20000-258
Apache HBase	2.6.3.7.3.2.20000-258
Apache Hive	3.1.3000.7.3.2.20000-258
Apache Iceberg	1.10.1.7.3.2.20000-258
Apache Impala	4.5.0.7.3.2.20000-258
Apache Kafka	3.9.2.7.3.2.20000-258
Apache Knox	2.1.0.7.3.2.20000-258
Apache Kudu	1.18.0.7.3.2.20000-258
Apache Livy	0.7.23000.7.3.2.20000-258
Apache Ozone	2.0.0.7.3.2.20000-258
Apache Oozie	5.1.0.7.3.2.20000-258
Apache ORC	2.0.3.7.3.2.20000-258
Apache Parquet	1.16.0.7.3.2.20000-258

Component	Version
Apache Phoenix	5.2.1.7.3.2.20000-258
Apache Ranger	2.8.0.7.3.2.20000-258
Apache Solr	8.11.2.7.3.2.20000-258
Apache Spark	3.5.8.7.3.2.20000-258
Apache Sqoop	1.4.7.7.3.2.20000-258
Apache Tez	0.9.1.7.3.2.20000-258
Apache ZooKeeper	3.8.5.7.3.2.20000-258

Other Components

Component	Version
Cruise Control	2.5.143.7.3.2.20000-258
Cloudera Lakehouse Optimizer (dlm)	1.0.0.7.3.2.20000-258
Cloudera Data Explorer (Hue)	4.11.0.7.3.2.20000-258
Search	1.0.0.7.3.2.20000-258
Schema Registry	3.5.8.7.3.2.20000-258
Streams Messaging Manager	2.3.0.7.3.2.20000-258
Streams Replication Manager	1.1.0.7.3.2.20000-258

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.3.2.20000-258
Hive Warehouse Connector	1.0.0.7.3.2.20000-258
Spark Atlas Connector	3.5.8.7.3.2.20000-258
Spark Schema Registry	3.5.8.7.3.2.20000-258

Cloudera Runtime 7.3.2.10000

List of the official component versions for Cloudera Runtime 7.3.2.10000.

Apache Components

Component	Version
Apache Arrow	15.0.0.7.3.2.10000-317
Apache Atlas	2.4.0.7.3.2.10000-317
Apache Calcite	1.33.0.7.3.2.10000-317
Apache Avatica	1.26.0.7.3.2.10000-317
Apache Avro	1.11.3.7.3.2.10000-317
Apache Hadoop (Includes HDFS, MapReduce and YARN)	3.4.2.7.3.2.10000-317
Apache HBase	2.6.3.7.3.2.10000-317
Apache Hive	3.1.3000.7.3.2.10000-317
Apache Iceberg	1.10.1.7.3.2.10000-317
Apache Impala	4.5.0.7.3.2.10000-317

Component	Version
Apache Kafka	3.9.2.7.3.2.10000-317
Apache Knox	2.1.0.7.3.2.10000-317
Apache Kudu	1.18.0.7.3.2.10000-317
Apache Livy	0.7.23000.7.3.2.10000-317
Apache Ozone	2.0.0.7.3.2.10000-317
Apache Oozie	5.1.0.7.3.2.10000-317
Apache ORC	2.0.3.7.3.2.10000-317
Apache Parquet	1.16.0.7.3.2.10000-317
Apache Phoenix	5.2.1.7.3.2.10000-317
Apache Ranger	2.8.0.7.3.2.10000-317
Apache Solr	8.11.2.7.3.2.10000-317
Apache Spark	3.5.8.7.3.2.10000-317
Apache Sqoop	1.4.7.7.3.2.10000-317
Apache Tez	0.9.1.7.3.2.10000-317
Apache ZooKeeper	3.8.5.7.3.2.10000-317

Other Components

Component	Version
Cruise Control	2.5.143.7.3.2.10000-317
Cloudera Lakehouse Optimizer (dlm)	1.0.0.7.3.2.10000-317
Cloudera Data Explorer (Hue)	4.11.0.7.3.2.10000-317
Search	1.0.0.7.3.2.10000-317
Schema Registry	3.5.8.7.3.2.10000-317
Streams Messaging Manager	2.3.0.7.3.2.10000-317
Streams Replication Manager	1.1.0.7.3.2.10000-317

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.3.2.10000-317
Hive Warehouse Connector	1.0.0.7.3.2.10000-317
Spark Atlas Connector	3.5.8.7.3.2.10000-317
Spark Schema Registry	3.5.8.7.3.2.10000-317

Cloudera Runtime 7.3.2.100

List of the official component versions for Cloudera Runtime 7.3.2.100.

Apache Components

Component	Version
Apache Arrow	15.0.0.7.3.2.100-296
Apache Atlas	2.4.0.7.3.2.100-296

Component	Version
Apache Calcite	1.33.0.7.3.2.100-296
Apache Avatica	1.26.0.7.3.2.100-296
Apache Avro	1.11.3.7.3.2.100-296
Apache Hadoop (Includes HDFS, MapReduce and YARN)	3.4.2.7.3.2.100-296
Apache HBase	2.6.3.7.3.2.100-296
Apache Hive	3.1.3000.7.3.2.100-296
Apache Iceberg	1.5.2.7.3.2.100-296
Apache Impala	4.5.0.7.3.2.100-296
Apache Kafka	3.9.1.7.3.2.100-296
Apache Knox	2.1.0.7.3.2.100-296
Apache Kudu	1.18.0.7.3.2.100-296
Apache Livy	0.7.23000.7.3.2.100-296
Apache Ozone	2.0.0.7.3.2.100-296
Apache Oozie	5.1.0.7.3.2.100-296
Apache ORC	2.0.3.7.3.2.100-296
Apache Parquet	1.15.2.7.3.2.100-296
Apache Phoenix	5.2.1.7.3.2.100-296
Apache Ranger	2.6.0.7.3.2.100-296
Apache Solr	8.11.2.7.3.2.100-296
Apache Spark	3.5.4.7.3.2.100-296
Apache Sqoop	1.4.7.7.3.2.100-296
Apache Tez	0.9.1.7.3.2.100-296
Apache ZooKeeper	3.8.5.7.3.2.100-296

Other Components

Component	Version
Cruise Control	2.5.143.7.3.2.100-296
Cloudera Lakehouse Optimizer (dlm)	1.0.0.7.3.2.100-296
Cloudera Data Explorer (Hue)	4.11.0.7.3.2.100-296
Search	1.0.0.7.3.2.100-296
Schema Registry	0.10.0.7.3.2.100-296
Streams Messaging Manager	2.3.0.7.3.2.100-296
Streams Replication Manager	1.1.0.7.3.2.100-296

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.3.2.100-296
Hive Warehouse Connector	1.0.0.7.3.2.100-296
Spark Atlas Connector	3.5.4.7.3.2.100-296
Spark Schema Registry	3.5.4.7.3.2.100-296

Cloudera Runtime 7.3.2

List of the official component versions for Cloudera Runtime 7.3.2.

Apache Components

Component	Version
Apache Arrow	15.0.0.7.3.2.0-957
Apache Atlas	2.4.0.7.3.2.0-957
Apache Calcite	1.33.0.7.3.2.0-957
Apache Avatica	1.26.0.7.3.2.0-957
Apache Avro	1.11.3.7.3.2.0-957
Apache Hadoop (Includes HDFS, MapReduce and YARN)	3.4.2.7.3.2.0-957
Apache HBase	2.6.3.7.3.2.0-957
Apache Hive	3.1.3000.7.3.2.0-957
Apache Iceberg	1.5.2.7.3.2.0-957
Apache Impala	4.5.0.7.3.2.0-957
Apache Kafka	3.9.1.7.3.2.0-957
Apache Knox	2.1.0.7.3.2.0-957
Apache Kudu	1.18.0.7.3.2.0-957
Apache Livy	0.7.23000.7.3.2.0-957
Apache Ozone	2.0.0.7.3.2.0-957
Apache Oozie	5.1.0.7.3.2.0-957
Apache ORC	2.0.3.7.3.2.0-957
Apache Parquet	1.15.2.7.3.2.0-957
Apache Phoenix	5.2.1.7.3.2.0-957
Apache Ranger	2.6.0.7.3.2.0-957
Apache Solr	8.11.2.7.3.2.0-957
Apache Spark	3.5.4.7.3.2.0-957
Apache Sqoop	1.4.7.7.3.2.0-957
Apache Tez	0.9.1.7.3.2.0-957
Apache ZooKeeper	3.8.5.7.3.2.0-957

Other Components

Component	Version
Cruise Control	2.5.143.7.3.2.0-957
Cloudera Lakehouse Optimizer (dlm)	1.0.0.7.3.2.0-957
Cloudera Data Explorer (Hue)	4.11.0.7.3.2.0-957
Search	1.0.0.7.3.2.0-957
Schema Registry	0.10.0.7.3.2.0-957
Streams Messaging Manager	2.3.0.7.3.2.0-957
Streams Replication Manager	1.1.0.7.3.2.0-957

Connectors and Encryption Components

Component	Version
HBase connectors	1.0.0.7.3.2.0-957
Hive Warehouse Connector	1.0.0.7.3.2.0-957
Spark Atlas Connector	3.5.4.7.3.2.0-957
Spark Schema Registry	3.5.4.7.3.2.0-957

Using the Cloudera Runtime Maven repository

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

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



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

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

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

Cloudera Runtime 7.3.2.20000-258

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

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-authn	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-authorization	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-classification-updater	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-client-common	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-client-v1	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-client-v2	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-common	2.4.0.7.3.2.20000-258

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-distro	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-docs	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-graphdb-api	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-graphdb-common	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-graphdb-janus	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-index-repair-tool	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-intg	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-notification	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-notification-analyzer	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-plugin-classloader	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-repository	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-server-api	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-server-common	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	atlas-testtools	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	hbase-bridge	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	hbase-bridge-shim	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	hbase-testing-util	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	hdfs-model	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	hive-bridge	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	hive-bridge-shim	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	impala-bridge	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	impala-bridge-shim	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	impala-hook-api	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	janusgraph-rdbms	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	kafka-bridge	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	kafka-bridge-shim	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	navigator-to-atlas	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	sample-app	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	sqoop-bridge	2.4.0.7.3.2.20000-258
Atlas	org.apache.atlas	sqoop-bridge-shim	2.4.0.7.3.2.20000-258
Avro	org.apache.avro	avro	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	avro-android	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	avro-codegen-test	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	avro-compiler	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	avro-grpc	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	avro-ipc	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	avro-ipc-jetty	1.11.3.7.3.2.20000-258

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-ipc-netty	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	avro-mapred	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	avro-maven-plugin	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	avro-perf	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	avro-protobuf	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	avro-service-archetype	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	avro-test-custom-conversions	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	avro-thrift	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	avro-tools	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	trevni-avro	1.11.3.7.3.2.20000-258
Avro	org.apache.avro	trevni-core	1.11.3.7.3.2.20000-258
Calcite	org.apache.calcite	calcite-babel	1.33.0.7.3.2.20000-258
Calcite	org.apache.calcite	calcite-core	1.33.0.7.3.2.20000-258
Calcite	org.apache.calcite	calcite-druid	1.33.0.7.3.2.20000-258
Calcite	org.apache.calcite	calcite-kafka	1.33.0.7.3.2.20000-258
Calcite	org.apache.calcite	calcite-linq4j	1.33.0.7.3.2.20000-258
Calcite	org.apache.calcite	calcite-plus	1.33.0.7.3.2.20000-258
Calcite	org.apache.calcite	calcite-server	1.33.0.7.3.2.20000-258
Calcite	org.apache.calcite	calcite-testkit	1.33.0.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-aliyun	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-annotations	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-archives	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-assemblies	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-auth	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-aws	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-azure	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-benchmark	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-build-tools	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-client	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-client-api	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-common	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-datajoin	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-distcp	3.4.2.7.3.2.20000-258

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-dynamometer-blockgen	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-dynamometer-dist	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-dynamometer-infra	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-dynamometer-workload	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-extras	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-federation-balance	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-fs2img	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-gridmix	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-hdfs	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-kafka	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-kms	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-minicluster	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-minikdc	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-nfs	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-openstack	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-registry	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-rumen	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-sls	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-streaming	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.4.2.7.3.2.20000-258

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-csi	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-globalpolicygenerator	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-documentstore	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.4.2.7.3.2.20000-258
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.4.2.7.3.2.20000-258
HBase	org.apache.hbase	hbase-annotations	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-asyncfs	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-backup	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-checkstyle	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-client	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-client-project	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-common	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-compression-aircompressor	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-compression-brotli	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-compression-lz4	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-compression-snappy	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-compression-zstd	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-endpoint	2.6.3.7.3.2.20000-258

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-examples	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-external-blockcache	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-hadoop-compat	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-hadoop2-compat	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-hbtop	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-http	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-it	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-logging	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-mapreduce	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-metrics	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-metrics-api	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-openssl	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-procedure	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-protocol	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-protocol-shaded	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-replication	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-resource-bundle	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-rest	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-rsgroup	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-server	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-shaded-client	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-shaded-client-project	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-shaded-testing-util	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-shell	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-testing-util	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-thrift	2.6.3.7.3.2.20000-258
HBase	org.apache.hbase	hbase-zookeeper	2.6.3.7.3.2.20000-258
Hive	org.apache.hive	catalogd-unit	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-beeline	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-classification	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-cli	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-common	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-contrib	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-exec	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.3.2.20000-258

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-hms-catalog	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-iceberg-catalog	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-iceberg-handler	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-iceberg-shading	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-impala	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-it-iceberg	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-it-iceberg-integration	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-it-util	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-jmh	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-metastore	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-parser	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-serde	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-service	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-shims	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-streaming	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-testutils	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-udf	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	kafka-handler	3.1.3000.7.3.2.20000-258

Project	groupId	artifactId	version
Hive	org.apache.hive	patched-iceberg-api	3.1.3000.7.3.2.20000-258
Hive	org.apache.hive	patched-iceberg-core	3.1.3000.7.3.2.20000-258
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector-spark3_2.12	1.0.0.7.3.2.20000-258
Kafka	org.apache.kafka	connect	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	connect-api	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	connect-basic-auth-extension	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	connect-cloudera-common	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	connect-file	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	connect-json	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	connect-mirror	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	connect-mirror-client	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	connect-runtime	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	connect-test-plugins	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	connect-transforms	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	generator	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-clients	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-examples	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-group-coordinator	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-group-coordinator-api	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-log4j-appender	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-metadata	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-raft	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-server	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-server-common	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-shell	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-storage	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-storage-api	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-examples	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-test-utils	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.9.2.7.3.2.20000-258

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-34	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-35	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-36	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-37	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-38	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-tools	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-tools-api	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka-transaction-coordinator	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka_2.12	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	kafka_2.13	3.9.2.7.3.2.20000-258
Kafka	org.apache.kafka	trogdor	3.9.2.7.3.2.20000-258
Knox	org.apache.knox	gateway-adapter	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-admin-ui	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-applications	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-cloud-bindings	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-demo-ldap	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-demo-ldap-launcher	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-discovery-ambari	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-discovery-cm	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-docker	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-i18n	2.1.0.7.3.2.20000-258

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-i18n-logging-log4j	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-i18n-logging-sl4j	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-openapi-ui	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-performance-test	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-ha	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-identity-assertion-common	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-identity-assertion-no-doas	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-jersey	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-rewrite	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-rewrite-common	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-security-authc-anon	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-security-authc-remote	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-security-authz-acls	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-security-authz-composite	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-security-authz-path-acls	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-security-clientcert	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-security-hadoopauth	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-security-jwt	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-security-pac4j	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-security-preauth	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-security-shiro	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-provider-security-webappsec	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-release	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-server	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-server-launcher	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-server-xforwarded-filter	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-admin	2.1.0.7.3.2.20000-258

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-as	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-auth	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-definitions	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-hashicorp-vault	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-hbase	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-health	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-hive	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-idbroker	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-idbroker-plugins	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-impala	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-jkg	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-knoxsso	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-knoxssout	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-knoxtoken	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-livy	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-metadata	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-nifi	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-nifi-registry	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-remoteconfig	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-restcatalog	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-rm	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-session	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-storm	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-test	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-tgs	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-vault	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-service-webhdfs	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-shell	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-shell-launcher	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-shell-release	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-shell-samples	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-spi	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-spi-common	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-test	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-test-idbroker	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-test-release-utils	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-test-utils	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-topology-hadoop-xml	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-topology-simple	2.1.0.7.3.2.20000-258

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-util-common	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-util-configinjector	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-util-launcher	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	gateway-util-urltemplate	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	hadoop-examples	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	knox-cli-launcher	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	knox-homepage-ui	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	knox-token-generation-ui	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	knox-token-management-ui	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	knox-webshell-ui	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	webhdfs-kerb-test	2.1.0.7.3.2.20000-258
Knox	org.apache.knox	webhdfs-test	2.1.0.7.3.2.20000-258
Kudu	org.apache.kudu	kudu-backup-common	1.18.0.7.3.2.20000-258
Kudu	org.apache.kudu	kudu-backup-tools	1.18.0.7.3.2.20000-258
Kudu	org.apache.kudu	kudu-backup3_2.12	1.18.0.7.3.2.20000-258
Kudu	org.apache.kudu	kudu-client	1.18.0.7.3.2.20000-258
Kudu	org.apache.kudu	kudu-flatbuffers	1.18.0.7.3.2.20000-258
Kudu	org.apache.kudu	kudu-hive	1.18.0.7.3.2.20000-258
Kudu	org.apache.kudu	kudu-proto	1.18.0.7.3.2.20000-258
Kudu	org.apache.kudu	kudu-replication	1.18.0.7.3.2.20000-258
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.18.0.7.3.2.20000-258
Kudu	org.apache.kudu	kudu-spark3_2.12	1.18.0.7.3.2.20000-258
Kudu	org.apache.kudu	kudu-subprocess	1.18.0.7.3.2.20000-258
Kudu	org.apache.kudu	kudu-test-utils	1.18.0.7.3.2.20000-258
Livy	org.apache.livy	livy-api	0.7.23000.7.3.2.20000-258
Livy	org.apache.livy	livy-client-common	0.7.23000.7.3.2.20000-258
Livy	org.apache.livy	livy-client-http	0.7.23000.7.3.2.20000-258
Livy	org.apache.livy	livy-core_2.12	0.7.23000.7.3.2.20000-258
Livy	org.apache.livy	livy-examples	0.7.23000.7.3.2.20000-258
Livy	org.apache.livy	livy-integration-test	0.7.23000.7.3.2.20000-258
Livy	org.apache.livy	livy-repl_2.12	0.7.23000.7.3.2.20000-258
Livy	org.apache.livy	livy-rsc	0.7.23000.7.3.2.20000-258
Livy	org.apache.livy	livy-scala-api_2.12	0.7.23000.7.3.2.20000-258
Livy	org.apache.livy	livy-server	0.7.23000.7.3.2.20000-258
Livy	org.apache.livy	livy-test-lib	0.7.23000.7.3.2.20000-258
Livy	org.apache.livy	livy-thriftserver	0.7.23000.7.3.2.20000-258
Livy	org.apache.livy	livy-thriftserver-session	0.7.23000.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.3.2.20000-258

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-analyzers-opennlp	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-core	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-join	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.3.2.20000-258
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-client	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-core	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-server	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.3.2.20000-258

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.3.2.20000-258
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.3.2.20000-258
ORC	org.apache.orc	orc-core	2.0.3.7.3.2.20000-258
ORC	org.apache.orc	orc-mapreduce	2.0.3.7.3.2.20000-258
ORC	org.apache.orc	orc-shims	2.0.3.7.3.2.20000-258
ORC	org.apache.orc	orc-tools	2.0.3.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-annotation-processing	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-cli-common	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-client	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-common	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-config	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-container-service	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-crypto-api	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-crypto-default	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-docs	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-erasurecode	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-interface-admin	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-interface-client	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-interface-server	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-managed-rocksdb	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-rocks-native	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-server-framework	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-server-scm	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	hdds-test-utils	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	mini-chaos-tests	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-cli-admin	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-cli-debug	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-cli-repair	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-cli-shell	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-client	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-common	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-csi	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-datanode	2.2.0.7.3.2.20000-258

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-dev-support	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-dist	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-filesystem	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-filesystem-common	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-filesystem-hadoop2	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-filesystem-hadoop3	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-filesystem-shaded	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-freon	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-httpfsgateway	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-iceberg	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-insight	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-integration-test	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-integration-test-recon	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-integration-test-s3	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-interface-client	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-interface-storage	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-manager	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-mini-cluster	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-multitenancy-ranger	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-network-tests	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-recon	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-reconcodegen	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-s3-secret-store	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-s3gateway	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-tools	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	ozone-vapor	2.2.0.7.3.2.20000-258
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	2.2.0.7.3.2.20000-258
Parquet	org.apache.parquet	parquet-avro	1.16.0.7.3.2.20000-258
Parquet	org.apache.parquet	parquet-column	1.16.0.7.3.2.20000-258
Parquet	org.apache.parquet	parquet-common	1.16.0.7.3.2.20000-258
Parquet	org.apache.parquet	parquet-encoding	1.16.0.7.3.2.20000-258
Parquet	org.apache.parquet	parquet-format-structures	1.16.0.7.3.2.20000-258
Parquet	org.apache.parquet	parquet-generator	1.16.0.7.3.2.20000-258
Parquet	org.apache.parquet	parquet-hadoop	1.16.0.7.3.2.20000-258
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.16.0.7.3.2.20000-258
Parquet	org.apache.parquet	parquet-jackson	1.16.0.7.3.2.20000-258
Parquet	org.apache.parquet	parquet-protobuf	1.16.0.7.3.2.20000-258
Parquet	org.apache.parquet	parquet-thrift	1.16.0.7.3.2.20000-258
Parquet	org.apache.parquet	parquet-tools	1.16.0.7.3.2.20000-258

Project	groupId	artifactId	version
Parquet	org.apache.parquet	parquet-variant	1.16.0.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.6	5.2.1.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-client-lite-hbase-2.6	5.2.1.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-core	5.2.1.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-core-client	5.2.1.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-core-server	5.2.1.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.2.1.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.0	5.2.1.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.4	5.2.1.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.6.0	5.2.1.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-mapreduce-byo-shaded-hbase-hbase-2.6	5.2.1.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-pherf	5.2.1.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.6	5.2.1.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.2.1.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix5-spark3	6.0.0.7.3.2.20000-258
Phoenix	org.apache.phoenix	phoenix5-spark3-shaded	6.0.0.7.3.2.20000-258
Ranger	org.apache.ranger	authz-embedded	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	authz-remote	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	conditions-enrichers	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	credentialbuilder	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	embeddedwebserver	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	hdfs-nn-perf	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	jisql	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ldapconfigcheck	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-adls-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-atlas-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-audit-core	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-audit-dest-cloudwatch	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-audit-dest-hdfs	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-audit-dest-kafka	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-audit-dest-log4j	2.8.0.7.3.2.20000-258

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-audit-dest-os	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-audit-dest-solr	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-audit-orc-util	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-authn	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-authz-api	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-caii-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-clo-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-common-ha	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-common-utils	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-distro	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-examples-distro	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-gs-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-hbase-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-hive-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-intg	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-kafka-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-kms	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-kms-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-knox-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-kudu-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-kylin-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-liquibase-upgrade	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-metrics	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-nestedstructure-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-nifi-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-opensearch-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-ozone-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-pdp	2.8.0.7.3.2.20000-258

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-plugin-classloader	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-plugin-perf	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-plugins-common	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-plugins-cred	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-plugins-installer	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-policymigration	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-raz-adls	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-raz-hook-s3	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-raz-intg	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-raz-processor	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-raz-s3	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-raz-s3-lib	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-rms-common	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-rms-hive	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-rms-tools	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-rms-webapp	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-s3-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-solr-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-storm-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-tagsync	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-tools	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-trino-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-util	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-yarn-plugin	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	sample-client	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	sampleapp	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	ugsync-util	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	unixauthclient	2.8.0.7.3.2.20000-258
Ranger	org.apache.ranger	unixauthservice	2.8.0.7.3.2.20000-258

Project	groupId	artifactId	version
Ranger	org.apache.ranger	unixusersync	2.8.0.7.3.2.20000-258
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-analytics	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-cell	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-core	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-langid	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-ltr	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-security-util	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-solrj	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-test-framework	8.11.2.7.3.2.20000-258
Solr	org.apache.solr	solr-velocity	8.11.2.7.3.2.20000-258
Spark	org.apache.spark	spark-avro_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-catalyst_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-common-utils_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-connect-client-jvm_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-connect-common_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-connect_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-core_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-graphx_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-hadoop-cloud_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-hive_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-kubernetes_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-kvstore_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-launcher_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-mllib-local_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-mllib_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-network-common_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-network-shuffle_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-network-yarn_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-protobuf_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-repl_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-shaded-raz	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-sketch_2.12	3.5.8.7.3.2.20000-258

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-sql-api_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-sql-kafka-0-10_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-sql_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-streaming_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-tags_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-unsafe_2.12	3.5.8.7.3.2.20000-258
Spark	org.apache.spark	spark-yarn_2.12	3.5.8.7.3.2.20000-258
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.3.2.20000-258
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.3.2.20000-258
Tez	org.apache.tez	hadoop-shim	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-api	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-aux-services	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-build-tools	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-common	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-dag	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-examples	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-history-parser	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-tests	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.3.2.20000-258
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.3.2.20000-258
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.5.7.3.2.20000-258
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.5.7.3.2.20000-258
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.5.7.3.2.20000-258
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.5.7.3.2.20000-258
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.5.7.3.2.20000-258
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.5.7.3.2.20000-258

Project	groupId	artifactId	version
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.5.7.3.2.20000-258
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.5.7.3.2.20000-258
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.5.7.3.2.20000-258
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.5.7.3.2.20000-258
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.5.7.3.2.20000-258

Cloudera Runtime 7.3.2.10000-317

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

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-authn	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-authorization	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-classification-updater	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-client-common	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-client-v1	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-client-v2	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-common	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-distro	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-docs	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-graphdb-api	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-graphdb-common	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-graphdb-janus	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-index-repair-tool	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-intg	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-janusgraph-hbase2	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-notification	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-notification-analyzer	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-plugin-classloader	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-repository	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-server-api	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-server-common	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	atlas-testtools	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	hbase-bridge	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	hbase-bridge-shim	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	hbase-testing-util	2.4.0.7.3.2.10000-317

Project	groupId	artifactId	version
Atlas	org.apache.atlas	hdfs-model	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	hive-bridge	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	hive-bridge-shim	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	impala-bridge	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	impala-bridge-shim	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	impala-hook-api	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	janusgraph-rdbms	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	kafka-bridge	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	kafka-bridge-shim	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	navigator-to-atlas	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	sample-app	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	sqoop-bridge	2.4.0.7.3.2.10000-317
Atlas	org.apache.atlas	sqoop-bridge-shim	2.4.0.7.3.2.10000-317
Avro	org.apache.avro	avro	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-android	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-codegen-test	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-compiler	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-grpc	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-ipc	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-ipc-jetty	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-ipc-netty	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-mapred	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-maven-plugin	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-perf	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-protobuf	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-service-archetype	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-test-custom-conversions	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-thrift	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	avro-tools	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	trevni-avro	1.11.3.7.3.2.10000-317
Avro	org.apache.avro	trevni-core	1.11.3.7.3.2.10000-317
Calcite	org.apache.calcite	calcite-babel	1.33.0.7.3.2.10000-317
Calcite	org.apache.calcite	calcite-core	1.33.0.7.3.2.10000-317
Calcite	org.apache.calcite	calcite-druid	1.33.0.7.3.2.10000-317
Calcite	org.apache.calcite	calcite-kafka	1.33.0.7.3.2.10000-317
Calcite	org.apache.calcite	calcite-linq4j	1.33.0.7.3.2.10000-317
Calcite	org.apache.calcite	calcite-plus	1.33.0.7.3.2.10000-317
Calcite	org.apache.calcite	calcite-server	1.33.0.7.3.2.10000-317
Calcite	org.apache.calcite	calcite-testkit	1.33.0.7.3.2.10000-317

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-aliyun	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-annotations	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-archives	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-assemblies	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-auth	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-aws	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-azure	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-benchmark	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-build-tools	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-client	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-client-api	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-common	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-datajoin	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-distcp	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-dynamometer-blockgen	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-dynamometer-dist	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-dynamometer-infra	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-dynamometer-workload	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-extras	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-federation-balance	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-fs2img	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-gridmix	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-hdfs	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-kafka	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-kms	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.4.2.7.3.2.10000-317

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-minicluster	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-minikdc	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-nfs	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-openstack	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-registry	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-rumen	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-sls	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-streaming	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-csi	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-globalpolicygenerator	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-documentstore	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.4.2.7.3.2.10000-317

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.4.2.7.3.2.10000-317
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.4.2.7.3.2.10000-317
HBase	org.apache.hbase	hbase-annotations	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-asyncfs	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-backup	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-checkstyle	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-client	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-client-project	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-common	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-compression-aircompressor	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-compression-brotli	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-compression-lz4	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-compression-snappy	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-compression-zstd	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-endpoint	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-examples	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-external-blockcache	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-hadoop-compat	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-hadoop2-compat	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-hbtop	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-http	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-it	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-logging	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-mapreduce	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-metrics	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-metrics-api	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-openssl	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-procedure	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-protocol	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-protocol-shaded	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-replication	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-resource-bundle	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-rest	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-rsgroup	2.6.3.7.3.2.10000-317

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-server	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-shaded-client	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-shaded-client-project	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-shaded-testing-util	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-shell	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-testing-util	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-thrift	2.6.3.7.3.2.10000-317
HBase	org.apache.hbase	hbase-zookeeper	2.6.3.7.3.2.10000-317
Hive	org.apache.hive	catalogd-unit	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-beeline	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-classification	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-cli	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-common	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-contrib	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-exec	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-hms-catalog	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-iceberg-catalog	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-iceberg-handler	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-iceberg-shading	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-impala	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-it-iceberg	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-it-iceberg-integration	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-it-util	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.3.2.10000-317

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-jmh	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-metastore	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-parser	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-serde	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-service	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-shims	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-streaming	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-testutils	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-udf	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	kafka-handler	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	patched-iceberg-api	3.1.3000.7.3.2.10000-317
Hive	org.apache.hive	patched-iceberg-core	3.1.3000.7.3.2.10000-317
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector-spark3_2.12	1.0.0.7.3.2.10000-317
Kafka	org.apache.kafka	connect	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	connect-api	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	connect-basic-auth-extension	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	connect-cloudera-common	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	connect-file	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	connect-json	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	connect-mirror	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	connect-mirror-client	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	connect-runtime	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	connect-test-plugins	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	connect-transforms	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	generator	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-clients	3.9.2.7.3.2.10000-317

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-examples	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-group-coordinator	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-group-coordinator-api	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-log4j-appender	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-metadata	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-raft	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-server	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-server-common	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-shell	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-storage	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-storage-api	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-examples	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-test-utils	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-34	3.9.2.7.3.2.10000-317

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-35	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-36	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-37	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-38	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-tools	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-tools-api	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka-transaction-coordinator	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka_2.12	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	kafka_2.13	3.9.2.7.3.2.10000-317
Kafka	org.apache.kafka	trogdor	3.9.2.7.3.2.10000-317
Knox	org.apache.knox	gateway-adapter	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-admin-ui	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-applications	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-cloud-bindings	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-demo-ldap	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-demo-ldap-launcher	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-discovery-ambari	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-discovery-cm	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-docker	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-i18n	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-i18n-logging-log4j	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-i18n-logging-sl4j	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-openapi-ui	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-performance-test	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-ha	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-identity-assertion-common	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-identity-assertion-no-doas	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-jersey	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-rewrite	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-rewrite-common	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	2.1.0.7.3.2.10000-317

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-security-authc-anon	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-security-authc-remote	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-security-authz-acls	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-security-authz-composite	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-security-authz-path-acls	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-security-clientcert	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-security-hadoopauth	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-security-jwt	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-security-pac4j	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-security-preauth	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-security-shiro	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-provider-security-webappsec	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-release	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-server	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-server-launcher	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-server-xforwarded-filter	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-admin	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-as	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-auth	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-definitions	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-hashicorp-vault	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-hbase	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-health	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-hive	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-idbroker	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-idbroker-plugins	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-impala	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-jkg	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-knoxsso	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-knoxssout	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-knoxtoken	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-livy	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-metadata	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-nifi	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-nifi-registry	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-remoteconfig	2.1.0.7.3.2.10000-317

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-restcatalog	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-rm	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-session	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-storm	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-test	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-tgs	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-vault	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-service-webhdfs	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-shell	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-shell-launcher	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-shell-release	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-shell-samples	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-spi	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-spi-common	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-test	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-test-idbroker	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-test-release-utils	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-test-utils	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-topology-hadoop-xml	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-topology-simple	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-util-common	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-util-configinjector	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-util-launcher	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	gateway-util-urltemplate	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	hadoop-examples	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	knox-cli-launcher	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	knox-homepage-ui	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	knox-token-generation-ui	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	knox-token-management-ui	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	knox-webshell-ui	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	webhdfs-kerb-test	2.1.0.7.3.2.10000-317
Knox	org.apache.knox	webhdfs-test	2.1.0.7.3.2.10000-317
Kudu	org.apache.kudu	kudu-backup-common	1.18.0.7.3.2.10000-317
Kudu	org.apache.kudu	kudu-backup-tools	1.18.0.7.3.2.10000-317
Kudu	org.apache.kudu	kudu-backup3_2.12	1.18.0.7.3.2.10000-317
Kudu	org.apache.kudu	kudu-client	1.18.0.7.3.2.10000-317
Kudu	org.apache.kudu	kudu-flatbuffers	1.18.0.7.3.2.10000-317
Kudu	org.apache.kudu	kudu-hive	1.18.0.7.3.2.10000-317
Kudu	org.apache.kudu	kudu-proto	1.18.0.7.3.2.10000-317

Project	groupId	artifactId	version
Kudu	org.apache.kudu	kudu-replication	1.18.0.7.3.2.10000-317
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.18.0.7.3.2.10000-317
Kudu	org.apache.kudu	kudu-spark3_2.12	1.18.0.7.3.2.10000-317
Kudu	org.apache.kudu	kudu-subprocess	1.18.0.7.3.2.10000-317
Kudu	org.apache.kudu	kudu-test-utils	1.18.0.7.3.2.10000-317
Livy	org.apache.livy	livy-api	0.7.23000.7.3.2.10000-317
Livy	org.apache.livy	livy-client-common	0.7.23000.7.3.2.10000-317
Livy	org.apache.livy	livy-client-http	0.7.23000.7.3.2.10000-317
Livy	org.apache.livy	livy-core_2.12	0.7.23000.7.3.2.10000-317
Livy	org.apache.livy	livy-examples	0.7.23000.7.3.2.10000-317
Livy	org.apache.livy	livy-integration-test	0.7.23000.7.3.2.10000-317
Livy	org.apache.livy	livy-repl_2.12	0.7.23000.7.3.2.10000-317
Livy	org.apache.livy	livy-rsc	0.7.23000.7.3.2.10000-317
Livy	org.apache.livy	livy-scala-api_2.12	0.7.23000.7.3.2.10000-317
Livy	org.apache.livy	livy-server	0.7.23000.7.3.2.10000-317
Livy	org.apache.livy	livy-test-lib	0.7.23000.7.3.2.10000-317
Livy	org.apache.livy	livy-thriftserver	0.7.23000.7.3.2.10000-317
Livy	org.apache.livy	livy-thriftserver-session	0.7.23000.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-analyzers-opennlp	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-core	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-join	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.3.2.10000-317

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.3.2.10000-317
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-client	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-core	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-server	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.3.2.10000-317
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.3.2.10000-317
ORC	org.apache.orc	orc-core	2.0.3.7.3.2.10000-317
ORC	org.apache.orc	orc-mapreduce	2.0.3.7.3.2.10000-317
ORC	org.apache.orc	orc-shims	2.0.3.7.3.2.10000-317
ORC	org.apache.orc	orc-tools	2.0.3.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-annotation-processing	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-cli-common	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-client	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-common	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-config	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-container-service	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-crypto-api	2.2.0.7.3.2.10000-317

Project	groupId	artifactId	version
Ozone	org.apache.ozone	hdds-crypto-default	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-docs	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-erasurecode	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-interface-admin	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-interface-client	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-interface-server	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-managed-rocksdb	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-rocks-native	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-server-framework	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-server-scm	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	hdds-test-utils	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	mini-chaos-tests	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-cli-admin	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-cli-debug	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-cli-repair	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-cli-shell	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-client	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-common	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-csi	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-datanode	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-dev-support	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-dist	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-filesystem	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-filesystem-common	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-filesystem-hadoop2	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-filesystem-hadoop3	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-filesystem-shaded	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-freon	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-httpfsgateway	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-iceberg	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-insight	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-integration-test	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-integration-test-recon	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-integration-test-s3	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-interface-client	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-interface-storage	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-manager	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-mini-cluster	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-multitenancy-ranger	2.2.0.7.3.2.10000-317

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-network-tests	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-recon	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-reconcodegen	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-s3-secret-store	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-s3gateway	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-tools	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	ozone-vapor	2.2.0.7.3.2.10000-317
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	2.2.0.7.3.2.10000-317
Parquet	org.apache.parquet	parquet-avro	1.16.0.7.3.2.10000-317
Parquet	org.apache.parquet	parquet-column	1.16.0.7.3.2.10000-317
Parquet	org.apache.parquet	parquet-common	1.16.0.7.3.2.10000-317
Parquet	org.apache.parquet	parquet-encoding	1.16.0.7.3.2.10000-317
Parquet	org.apache.parquet	parquet-format-structures	1.16.0.7.3.2.10000-317
Parquet	org.apache.parquet	parquet-generator	1.16.0.7.3.2.10000-317
Parquet	org.apache.parquet	parquet-hadoop	1.16.0.7.3.2.10000-317
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.16.0.7.3.2.10000-317
Parquet	org.apache.parquet	parquet-jackson	1.16.0.7.3.2.10000-317
Parquet	org.apache.parquet	parquet-protobuf	1.16.0.7.3.2.10000-317
Parquet	org.apache.parquet	parquet-thrift	1.16.0.7.3.2.10000-317
Parquet	org.apache.parquet	parquet-tools	1.16.0.7.3.2.10000-317
Parquet	org.apache.parquet	parquet-variant	1.16.0.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.6	5.2.1.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-client-lite-hbase-2.6	5.2.1.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-core	5.2.1.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-core-client	5.2.1.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-core-server	5.2.1.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.2.1.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.0	5.2.1.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.4	5.2.1.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.6.0	5.2.1.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-mapreduce-byo-shaded-hbase-hbase-2.6	5.2.1.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-pherf	5.2.1.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.6	5.2.1.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.2.1.7.3.2.10000-317

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix5-spark3	6.0.0.7.3.2.10000-317
Phoenix	org.apache.phoenix	phoenix5-spark3-shaded	6.0.0.7.3.2.10000-317
Ranger	org.apache.ranger	authz-embedded	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	authz-remote	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	conditions-enrichers	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	credentialbuilder	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	embeddedwebserver	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	hdfs-nn-perf	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	jisql	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ldapconfigcheck	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-adls-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-atlas-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-audit-core	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-audit-dest-cloudwatch	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-audit-dest-hdfs	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-audit-dest-kafka	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-audit-dest-log4j	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-audit-dest-os	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-audit-dest-solr	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-audit-orc-util	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-authn	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-authz-api	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-caii-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-clo-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-common-ha	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-common-utils	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-distro	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-examples-distro	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-gs-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-hbase-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-hive-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-intg	2.8.0.7.3.2.10000-317

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-kafka-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-kms	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-kms-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-knox-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-kudu-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-kylin-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-liquibase-upgrade	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-metrics	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-nestedstructure-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-nifi-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-opensearch-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-ozone-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-pdp	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-plugin-classloader	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-plugin-perf	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-plugins-common	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-plugins-cred	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-plugins-installer	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-policymigration	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-raz-adls	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-raz-hook-s3	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-raz-intg	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-raz-processor	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-raz-s3	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-raz-s3-lib	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-rms-common	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-rms-hive	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-rms-tools	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-rms-webapp	2.8.0.7.3.2.10000-317

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-s3-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-solr-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-storm-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-tagsync	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-tools	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-trino-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-util	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-yarn-plugin	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	sample-client	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	sampleapp	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	ugsync-util	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	unixauthclient	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	unixauthservice	2.8.0.7.3.2.10000-317
Ranger	org.apache.ranger	unixusersync	2.8.0.7.3.2.10000-317
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-analytics	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-cell	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-core	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-langid	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-ltr	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-security-util	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-solrj	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-test-framework	8.11.2.7.3.2.10000-317
Solr	org.apache.solr	solr-velocity	8.11.2.7.3.2.10000-317
Spark	org.apache.spark	spark-avro_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-catalyst_2.12	3.5.8.7.3.2.10000-317

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-common-utils_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-connect-client-jvm_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-connect-common_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-connect_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-core_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-graphx_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-hadoop-cloud_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-hive_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-kubernetes_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-kvstore_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-launcher_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-mllib-local_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-mllib_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-network-common_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-network-shuffle_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-network-yarn_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-protobuf_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-repl_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-shaded-raz	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-sketch_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-sql-api_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-sql-kafka-0-10_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-sql_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-streaming_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-tags_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-unsafe_2.12	3.5.8.7.3.2.10000-317
Spark	org.apache.spark	spark-yarn_2.12	3.5.8.7.3.2.10000-317
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.3.2.10000-317
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.3.2.10000-317
Tez	org.apache.tez	hadoop-shim	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-api	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-aux-services	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-build-tools	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-common	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-dag	0.9.1.7.3.2.10000-317

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-examples	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-history-parser	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-tests	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.3.2.10000-317
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.3.2.10000-317
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.5.7.3.2.10000-317
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.5.7.3.2.10000-317
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.5.7.3.2.10000-317
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.5.7.3.2.10000-317
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.5.7.3.2.10000-317
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.5.7.3.2.10000-317
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.5.7.3.2.10000-317
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.5.7.3.2.10000-317
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.5.7.3.2.10000-317
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.5.7.3.2.10000-317
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.5.7.3.2.10000-317

Cloudera Runtime 7.3.2.100-296

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

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-authorization	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-classification-updater	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-client-common	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-client-v1	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-client-v2	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-common	2.4.0.7.3.2.100-296

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-distro	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-docs	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-graphdb-api	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-graphdb-common	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-graphdb-janus	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-index-repair-tool	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-intg	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-notification	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-notification-analyzer	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-plugin-classloader	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-repository	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-server-api	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	atlas-testtools	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	hbase-bridge	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	hbase-bridge-shim	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	hbase-testing-util	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	hdfs-model	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	hive-bridge	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	hive-bridge-shim	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	impala-bridge	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	impala-bridge-shim	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	impala-hook-api	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	kafka-bridge	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	kafka-bridge-shim	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	navigator-to-atlas	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	sample-app	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	sqoop-bridge	2.4.0.7.3.2.100-296
Atlas	org.apache.atlas	sqoop-bridge-shim	2.4.0.7.3.2.100-296
Avro	org.apache.avro	avro	1.11.3.7.3.2.100-296
Avro	org.apache.avro	avro-android	1.11.3.7.3.2.100-296
Avro	org.apache.avro	avro-codegen-test	1.11.3.7.3.2.100-296
Avro	org.apache.avro	avro-compiler	1.11.3.7.3.2.100-296
Avro	org.apache.avro	avro-grpc	1.11.3.7.3.2.100-296
Avro	org.apache.avro	avro-ipc	1.11.3.7.3.2.100-296
Avro	org.apache.avro	avro-ipc-jetty	1.11.3.7.3.2.100-296
Avro	org.apache.avro	avro-ipc-netty	1.11.3.7.3.2.100-296
Avro	org.apache.avro	avro-mapred	1.11.3.7.3.2.100-296
Avro	org.apache.avro	avro-maven-plugin	1.11.3.7.3.2.100-296

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-perf	1.11.3.7.3.2.100-296
Avro	org.apache.avro	avro-protobuf	1.11.3.7.3.2.100-296
Avro	org.apache.avro	avro-service-archetype	1.11.3.7.3.2.100-296
Avro	org.apache.avro	avro-test-custom-conversions	1.11.3.7.3.2.100-296
Avro	org.apache.avro	avro-thrift	1.11.3.7.3.2.100-296
Avro	org.apache.avro	avro-tools	1.11.3.7.3.2.100-296
Avro	org.apache.avro	trevni-avro	1.11.3.7.3.2.100-296
Avro	org.apache.avro	trevni-core	1.11.3.7.3.2.100-296
Calcite	org.apache.calcite	calcite-babel	1.33.0.7.3.2.100-296
Calcite	org.apache.calcite	calcite-core	1.33.0.7.3.2.100-296
Calcite	org.apache.calcite	calcite-druid	1.33.0.7.3.2.100-296
Calcite	org.apache.calcite	calcite-kafka	1.33.0.7.3.2.100-296
Calcite	org.apache.calcite	calcite-linq4j	1.33.0.7.3.2.100-296
Calcite	org.apache.calcite	calcite-plus	1.33.0.7.3.2.100-296
Calcite	org.apache.calcite	calcite-server	1.33.0.7.3.2.100-296
Calcite	org.apache.calcite	calcite-testkit	1.33.0.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-aliyun	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-annotations	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-archives	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-assemblies	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-auth	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-aws	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-azure	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-benchmark	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-build-tools	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-client	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-client-api	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-common	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-datajoin	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-distcp	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-dynamometer-blockgen	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-dynamometer-dist	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-dynamometer-infra	3.4.2.7.3.2.100-296

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-dynamometer-workload	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-extras	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-federation-balance	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-fs2img	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-gridmix	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-hdfs	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-kafka	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-kms	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-minicluster	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-minikdc	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-nfs	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-openstack	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-registry	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-rumen	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-sls	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-streaming	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.4.2.7.3.2.100-296

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-csi	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-globalpolicygenerator	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-documentstore	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.4.2.7.3.2.100-296
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.4.2.7.3.2.100-296
HBase	org.apache.hbase	hbase-annotations	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-asyncfs	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-backup	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-checkstyle	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-client	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-client-project	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-common	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-compression-aircompressor	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-compression-brotli	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-compression-lz4	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-compression-snappy	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-compression-zstd	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-endpoint	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-examples	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-external-blockcache	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-hadoop-compat	2.6.3.7.3.2.100-296

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-hadoop2-compat	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-hbtop	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-http	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-it	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-logging	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-mapreduce	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-metrics	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-metrics-api	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-openssl	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-procedure	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-protocol	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-protocol-shaded	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-replication	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-resource-bundle	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-rest	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-rsgroup	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-server	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-shaded-client	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-shaded-client-project	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-shaded-testing-util	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-shell	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-testing-util	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-thrift	2.6.3.7.3.2.100-296
HBase	org.apache.hbase	hbase-zookeeper	2.6.3.7.3.2.100-296
Hive	org.apache.hive	catalogd-unit	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-beeline	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-classification	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-cli	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-common	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-contrib	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-exec	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-hms-catalog	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.3.2.100-296

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-iceberg-catalog	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-iceberg-handler	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-iceberg-shading	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-impala	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-it-iceberg	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-it-impala	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-it-util	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-jmh	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-metastore	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-parser	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-serde	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-service	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-shims	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-streaming	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-testutils	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-udf	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	kafka-handler	3.1.3000.7.3.2.100-296
Hive	org.apache.hive	patched-iceberg-api	patched-1.5.2.7.3.2.100-296-3.1.3000.7.3.2.100-2
Hive	org.apache.hive	patched-iceberg-core	patched-1.5.2.7.3.2.100-296-3.1.3000.7.3.2.100-2
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector-spark3_2.12	1.0.0.7.3.2.100-296

Project	groupId	artifactId	version
Kafka	org.apache.kafka	connect	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	connect-api	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	connect-basic-auth-extension	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	connect-cloudera-common	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	connect-file	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	connect-json	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	connect-mirror	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	connect-mirror-client	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	connect-runtime	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	connect-test-plugins	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	connect-transforms	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	generator	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-clients	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-examples	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-group-coordinator	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-group-coordinator-api	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-log4j-appender	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-metadata	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-raft	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-server	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-server-common	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-shell	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-storage	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-storage-api	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-examples	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-test-utils	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.9.1.7.3.2.100-296

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-34	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-35	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-36	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-37	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-38	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-tools	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-tools-api	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka-transaction-coordinator	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka_2.12	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	kafka_2.13	3.9.1.7.3.2.100-296
Kafka	org.apache.kafka	trogdor	3.9.1.7.3.2.100-296
Knox	org.apache.knox	gateway-adapter	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-admin-ui	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-applications	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-cloud-bindings	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-demo-ldap	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-demo-ldap-launcher	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-discovery-ambari	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-discovery-cm	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-docker	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-i18n	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-i18n-logging-log4j	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-i18n-logging-sl4j	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-openapi-ui	2.1.0.7.3.2.100-296

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-performance-test	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-ha	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-identity-assertion-common	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-identity-assertion-no-doas	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-jersey	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-rewrite	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-rewrite-common	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-security-authc-anon	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-security-authc-remote	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-security-authz-acls	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-security-authz-composite	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-security-authz-path-acls	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-security-clientcert	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-security-hadoopauth	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-security-jwt	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-security-pac4j	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-security-preauth	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-security-shiro	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-provider-security-webappsec	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-release	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-server	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-server-launcher	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-server-xforwarded-filter	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-admin	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-as	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-auth	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-definitions	2.1.0.7.3.2.100-296

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-hashicorp-vault	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-hbase	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-health	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-hive	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-idbroker	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-idbroker-plugins	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-impala	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-jkg	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-knoxsso	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-knoxssout	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-knoxtoken	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-livy	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-metadata	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-nifi	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-nifi-registry	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-remoteconfig	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-rm	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-session	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-storm	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-test	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-tgs	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-vault	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-service-webhdfs	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-shell	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-shell-launcher	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-shell-release	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-shell-samples	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-spi	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-spi-common	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-test	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-test-idbroker	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-test-release-utils	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-test-utils	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-topology-hadoop-xml	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-topology-simple	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-util-common	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-util-configinjector	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-util-launcher	2.1.0.7.3.2.100-296
Knox	org.apache.knox	gateway-util-urltemplate	2.1.0.7.3.2.100-296

Project	groupId	artifactId	version
Knox	org.apache.knox	hadoop-examples	2.1.0.7.3.2.100-296
Knox	org.apache.knox	knox-cli-launcher	2.1.0.7.3.2.100-296
Knox	org.apache.knox	knox-homepage-ui	2.1.0.7.3.2.100-296
Knox	org.apache.knox	knox-token-generation-ui	2.1.0.7.3.2.100-296
Knox	org.apache.knox	knox-token-management-ui	2.1.0.7.3.2.100-296
Knox	org.apache.knox	knox-webshell-ui	2.1.0.7.3.2.100-296
Knox	org.apache.knox	webhdfs-kerb-test	2.1.0.7.3.2.100-296
Knox	org.apache.knox	webhdfs-test	2.1.0.7.3.2.100-296
Kudu	org.apache.kudu	kudu-backup-common	1.18.0.7.3.2.100-296
Kudu	org.apache.kudu	kudu-backup-tools	1.18.0.7.3.2.100-296
Kudu	org.apache.kudu	kudu-backup3_2.12	1.18.0.7.3.2.100-296
Kudu	org.apache.kudu	kudu-client	1.18.0.7.3.2.100-296
Kudu	org.apache.kudu	kudu-flatbuffers	1.18.0.7.3.2.100-296
Kudu	org.apache.kudu	kudu-hive	1.18.0.7.3.2.100-296
Kudu	org.apache.kudu	kudu-proto	1.18.0.7.3.2.100-296
Kudu	org.apache.kudu	kudu-replication	1.18.0.7.3.2.100-296
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.18.0.7.3.2.100-296
Kudu	org.apache.kudu	kudu-spark3_2.12	1.18.0.7.3.2.100-296
Kudu	org.apache.kudu	kudu-subprocess	1.18.0.7.3.2.100-296
Kudu	org.apache.kudu	kudu-test-utils	1.18.0.7.3.2.100-296
Livy	org.apache.livy	livy-api	0.7.23000.7.3.2.100-296
Livy	org.apache.livy	livy-client-common	0.7.23000.7.3.2.100-296
Livy	org.apache.livy	livy-client-http	0.7.23000.7.3.2.100-296
Livy	org.apache.livy	livy-core_2.12	0.7.23000.7.3.2.100-296
Livy	org.apache.livy	livy-examples	0.7.23000.7.3.2.100-296
Livy	org.apache.livy	livy-integration-test	0.7.23000.7.3.2.100-296
Livy	org.apache.livy	livy-repl_2.12	0.7.23000.7.3.2.100-296
Livy	org.apache.livy	livy-rsc	0.7.23000.7.3.2.100-296
Livy	org.apache.livy	livy-scala-api_2.12	0.7.23000.7.3.2.100-296
Livy	org.apache.livy	livy-server	0.7.23000.7.3.2.100-296
Livy	org.apache.livy	livy-test-lib	0.7.23000.7.3.2.100-296
Livy	org.apache.livy	livy-thriftserver	0.7.23000.7.3.2.100-296
Livy	org.apache.livy	livy-thriftserver-session	0.7.23000.7.3.2.100-296
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-analyzers-opennlp	8.11.2.7.3.2.100-296

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-core	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-join	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.3.2.100-296
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.3.2.100-296
Oozie	org.apache.oozie	oozie-client	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-core	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-server	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.3.2.100-296

Project	groupId	artifactId	version
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.3.2.100-296
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.3.2.100-296
ORC	org.apache.orc	orc-core	2.0.3.7.3.2.100-296
ORC	org.apache.orc	orc-mapreduce	2.0.3.7.3.2.100-296
ORC	org.apache.orc	orc-shims	2.0.3.7.3.2.100-296
ORC	org.apache.orc	orc-tools	2.0.3.7.3.2.100-296
Ozone	org.apache.ozone	hdds-annotation-processing	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-client	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-common	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-config	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-container-service	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-crypto-api	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-crypto-default	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-docs	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-erasurecode	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-interface-admin	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-interface-client	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-interface-server	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-managed-rocksdb	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-rocks-native	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-server-framework	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-server-scm	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	hdds-test-utils	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	mini-chaos-tests	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-cli-admin	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-cli-shell	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-client	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-common	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-csi	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-datanode	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-dev-support	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-dist	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-filesystem	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-filesystem-common	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-filesystem-hadoop2	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-filesystem-hadoop3	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-filesystem-shaded	2.0.0.7.3.2.100-296

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-freon	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-httpfsgateway	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-insight	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-integration-test	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-integration-test-recon	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-integration-test-s3	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-interface-client	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-interface-storage	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-manager	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-mini-cluster	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-network-tests	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-recon	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-reconcodegen	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-s3-secret-store	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-s3gateway	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	ozone-tools	2.0.0.7.3.2.100-296
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	2.0.0.7.3.2.100-296
Parquet	org.apache.parquet	parquet-avro	1.15.2.7.3.2.100-296
Parquet	org.apache.parquet	parquet-column	1.15.2.7.3.2.100-296
Parquet	org.apache.parquet	parquet-common	1.15.2.7.3.2.100-296
Parquet	org.apache.parquet	parquet-encoding	1.15.2.7.3.2.100-296
Parquet	org.apache.parquet	parquet-format-structures	1.15.2.7.3.2.100-296
Parquet	org.apache.parquet	parquet-generator	1.15.2.7.3.2.100-296
Parquet	org.apache.parquet	parquet-hadoop	1.15.2.7.3.2.100-296
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.15.2.7.3.2.100-296
Parquet	org.apache.parquet	parquet-jackson	1.15.2.7.3.2.100-296
Parquet	org.apache.parquet	parquet-pig	1.15.2.7.3.2.100-296
Parquet	org.apache.parquet	parquet-pig-bundle	1.15.2.7.3.2.100-296
Parquet	org.apache.parquet	parquet-protobuf	1.15.2.7.3.2.100-296
Parquet	org.apache.parquet	parquet-thrift	1.15.2.7.3.2.100-296
Parquet	org.apache.parquet	parquet-tools	1.15.2.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.6	5.2.1.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-client-lite-hbase-2.6	5.2.1.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-core	5.2.1.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-core-client	5.2.1.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-core-server	5.2.1.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.2.1.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.0	5.2.1.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.4	5.2.1.7.3.2.100-296

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.6.0	5.2.1.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-mapreduce-byo-shaded-hbase-hbase-2.6	5.2.1.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-pherf	5.2.1.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.6	5.2.1.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.2.1.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix5-spark3	6.0.0.7.3.2.100-296
Phoenix	org.apache.phoenix	phoenix5-spark3-shaded	6.0.0.7.3.2.100-296
Ranger	org.apache.ranger	conditions-enrichers	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	credentialbuilder	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	embeddedwebserver	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	hdfs-nn-perf	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	jisql	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ldapconfigcheck	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-adls-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-atlas-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-authn	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-caii-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-clo-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-common-ha	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-distro	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-examples-distro	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-gs-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-hbase-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-hive-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-intg	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-kafka-plugin	2.6.0.7.3.2.100-296

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-kms	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-kms-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-knox-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-kudu-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-kylin-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-metrics	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-nestedstructure-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-nifi-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-opensearch-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-ozone-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-plugin-classloader	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-plugin-perf	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-plugins-audit	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-plugins-common	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-plugins-cred	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-plugins-installer	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-policymigration	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-raz-adls	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-raz-hook-s3	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-raz-intg	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-raz-processor	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-raz-s3	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-raz-s3-lib	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-rms-common	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-rms-hive	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-rms-tools	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-rms-webapp	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-s3-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.6.0.7.3.2.100-296

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-solr-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-storm-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-tagsync	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-tools	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-trino-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-util	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-yarn-plugin	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	sample-client	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	sampleapp	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	ugsync-util	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	unixauthclient	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	unixauthservice	2.6.0.7.3.2.100-296
Ranger	org.apache.ranger	unixusersync	2.6.0.7.3.2.100-296
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-analytics	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-cell	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-core	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-langid	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-ltr	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-security-util	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-solrj	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-test-framework	8.11.2.7.3.2.100-296
Solr	org.apache.solr	solr-velocity	8.11.2.7.3.2.100-296
Spark	org.apache.spark	spark-avro_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-catalyst_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-common-utils_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-connect-client-jvm_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-connect-common_2.12	3.5.4.7.3.2.100-296

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-connect_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-core_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-graphx_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-hadoop-cloud_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-hive_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-kubernetes_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-kvstore_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-launcher_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-mllib-local_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-mllib_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-network-common_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-network-shuffle_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-network-yarn_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-protobuf_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-repl_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-shaded-raz	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-sketch_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-sql-api_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-sql-kafka-0-10_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-sql_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-streaming_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-tags_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-unsafe_2.12	3.5.4.7.3.2.100-296
Spark	org.apache.spark	spark-yarn_2.12	3.5.4.7.3.2.100-296
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.3.2.100-296
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.3.2.100-296
Tez	org.apache.tez	hadoop-shim	0.9.1.7.3.2.100-296
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-api	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-aux-services	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-build-tools	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-common	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-dag	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-examples	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-history-parser	0.9.1.7.3.2.100-296

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-tests	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.3.2.100-296
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.3.2.100-296
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.5.7.3.2.100-296
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.5.7.3.2.100-296
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.5.7.3.2.100-296
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.5.7.3.2.100-296
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.5.7.3.2.100-296
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.5.7.3.2.100-296
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.5.7.3.2.100-296
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.5.7.3.2.100-296
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.5.7.3.2.100-296
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.5.7.3.2.100-296
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.5.7.3.2.100-296

Cloudera Runtime 7.3.2.0-957

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

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-authorization	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-aws-s3-bridge	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-azure-adls-bridge	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-classification-updater	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-client-common	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-client-v1	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-client-v2	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-client-v2-shaded	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-common	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-distro	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-docs	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-graphdb-api	2.4.0.7.3.2.0-957

Project	groupId	artifactId	version
Atlas	org.apache.atlas	atlas-graphdb-common	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-graphdb-janus	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-hdfs-bridge	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-index-repair-tool	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-intg	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-notification	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-notification-analyzer	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-plugin-classloader	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-repository	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-server-api	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	atlas-testtools	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	hbase-bridge	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	hbase-bridge-shim	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	hbase-testing-util	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	hdfs-model	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	hive-bridge	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	hive-bridge-shim	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	impala-bridge	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	impala-bridge-shim	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	impala-hook-api	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	kafka-bridge	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	kafka-bridge-shim	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	navigator-to-atlas	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	sample-app	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	sqoop-bridge	2.4.0.7.3.2.0-957
Atlas	org.apache.atlas	sqoop-bridge-shim	2.4.0.7.3.2.0-957
Avro	org.apache.avro	avro	1.11.3.7.3.2.0-957
Avro	org.apache.avro	avro-android	1.11.3.7.3.2.0-957
Avro	org.apache.avro	avro-codegen-test	1.11.3.7.3.2.0-957
Avro	org.apache.avro	avro-compiler	1.11.3.7.3.2.0-957
Avro	org.apache.avro	avro-grpc	1.11.3.7.3.2.0-957
Avro	org.apache.avro	avro-ipc	1.11.3.7.3.2.0-957
Avro	org.apache.avro	avro-ipc-jetty	1.11.3.7.3.2.0-957
Avro	org.apache.avro	avro-ipc-netty	1.11.3.7.3.2.0-957
Avro	org.apache.avro	avro-mapred	1.11.3.7.3.2.0-957
Avro	org.apache.avro	avro-maven-plugin	1.11.3.7.3.2.0-957
Avro	org.apache.avro	avro-perf	1.11.3.7.3.2.0-957
Avro	org.apache.avro	avro-protobuf	1.11.3.7.3.2.0-957
Avro	org.apache.avro	avro-service-archetype	1.11.3.7.3.2.0-957

Project	groupId	artifactId	version
Avro	org.apache.avro	avro-test-custom-conversions	1.11.3.7.3.2.0-957
Avro	org.apache.avro	avro-thrift	1.11.3.7.3.2.0-957
Avro	org.apache.avro	avro-tools	1.11.3.7.3.2.0-957
Avro	org.apache.avro	trevni-avro	1.11.3.7.3.2.0-957
Avro	org.apache.avro	trevni-core	1.11.3.7.3.2.0-957
Calcite	org.apache.calcite	calcite-babel	1.33.0.7.3.2.0-957
Calcite	org.apache.calcite	calcite-core	1.33.0.7.3.2.0-957
Calcite	org.apache.calcite	calcite-druid	1.33.0.7.3.2.0-957
Calcite	org.apache.calcite	calcite-kafka	1.33.0.7.3.2.0-957
Calcite	org.apache.calcite	calcite-linq4j	1.33.0.7.3.2.0-957
Calcite	org.apache.calcite	calcite-plus	1.33.0.7.3.2.0-957
Calcite	org.apache.calcite	calcite-server	1.33.0.7.3.2.0-957
Calcite	org.apache.calcite	calcite-testkit	1.33.0.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-aliyun	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-annotations	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-archive-logs	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-archives	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-assemblies	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-auth	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-aws	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-azure	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-azure-datalake	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-benchmark	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-build-tools	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-client	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-client-api	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-client-integration-tests	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-client-minicluster	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-client-runtime	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-cloud-storage	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-common	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-datajoin	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-distcp	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-dynamometer-blockgen	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-dynamometer-dist	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-dynamometer-infra	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-dynamometer-workload	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-extras	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-federation-balance	3.4.2.7.3.2.0-957

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-fs2img	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-gridmix	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-hdfs	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-hdfs-client	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-hdfs-httpfs	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-hdfs-native-client	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-hdfs-nfs	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-hdfs-rbf	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-kafka	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-kms	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-app	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-common	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-core	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-hs-plugins	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-jobclient	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-nativetask	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-shuffle	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-mapreduce-client-uploader	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-mapreduce-examples	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-maven-plugins	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-minicluster	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-minikdc	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-nfs	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-openstack	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-registry	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-resourceestimator	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-rumen	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-sls	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-streaming	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-tools-dist	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-api	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-applications-distributedshell	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-applications-unmanaged-am-launcher	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-client	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-common	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-csi	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-registry	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-applicationhistoryservice	3.4.2.7.3.2.0-957

Project	groupId	artifactId	version
Hadoop	org.apache.hadoop	hadoop-yarn-server-common	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-globalpolicygenerator	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-nodemanager	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-resourcemanager	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-router	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-sharedcachemanager	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-tests	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-timeline-pluginstorage	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-documentstore	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-client	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-common	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-server-2	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-timelineservice-hbase-tests	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-server-web-proxy	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-services-api	3.4.2.7.3.2.0-957
Hadoop	org.apache.hadoop	hadoop-yarn-services-core	3.4.2.7.3.2.0-957
HBase	org.apache.hbase	hbase-annotations	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-asyncfs	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-backup	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-checkstyle	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-client	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-client-project	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-common	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-compression-aircompressor	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-compression-brotli	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-compression-lz4	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-compression-snappy	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-compression-zstd	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-endpoint	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-examples	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-external-blockcache	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-hadoop-compat	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-hadoop2-compat	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-hbtop	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-http	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-it	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-logging	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-mapreduce	2.6.3.7.3.2.0-957

Project	groupId	artifactId	version
HBase	org.apache.hbase	hbase-metrics	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-metrics-api	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-openssl	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-procedure	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-protocol	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-protocol-shaded	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-replication	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-resource-bundle	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-rest	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-rsgroup	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-server	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-shaded-client	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-shaded-client-byo-hadoop	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-shaded-client-project	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-shaded-mapreduce	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-shaded-testing-util	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-shaded-testing-util-tester	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-shell	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-testing-util	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-thrift	2.6.3.7.3.2.0-957
HBase	org.apache.hbase	hbase-zookeeper	2.6.3.7.3.2.0-957
Hive	org.apache.hive	catalogd-unit	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-beeline	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-blobstore	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-classification	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-cli	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-common	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-contrib	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-exec	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-hbase-handler	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-hcatalog-it-unit	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-hms-catalog	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-hplsql	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-iceberg-catalog	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-iceberg-handler	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-iceberg-shading	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-impala	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-it-custom-serde	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-it-iceberg	3.1.3000.7.3.2.0-957

Project	groupId	artifactId	version
Hive	org.apache.hive	hive-it-impala	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-it-minikdc	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-it-qfile	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-it-qfile-kudu	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-it-test-serde	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-it-unit	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-it-unit-hadoop2	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-it-util	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-jdbc	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-jdbc-handler	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-jmh	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-kudu-handler	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-llap-client	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-llap-common	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-llap-ext-client	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-llap-server	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-llap-tez	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-metastore	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-parser	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-serde	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-service	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-service-rpc	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-shims	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-standalone-metastore	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-storage-api	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-streaming	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-testutils	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-udf	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	hive-vector-code-gen	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	kafka-handler	3.1.3000.7.3.2.0-957
Hive	org.apache.hive	patched-iceberg-api	patched-1.5.2.7.3.2.0-957-3.1.3000.7
Hive	org.apache.hive	patched-iceberg-core	patched-1.5.2.7.3.2.0-957-3.1.3000.7
Hive Warehouse Connector	com.hortonworks.hive	hive-warehouse-connector-spark3_2.12	1.0.0.7.3.2.0-957
Kafka	org.apache.kafka	connect	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	connect-api	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	connect-basic-auth-extension	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	connect-cloudera-authorization-extension	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	connect-cloudera-common	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	connect-cloudera-secret-storage	3.9.1.7.3.2.0-957

Project	groupId	artifactId	version
Kafka	org.apache.kafka	connect-cloudera-security-policies	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	connect-file	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	connect-json	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	connect-mirror	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	connect-mirror-client	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	connect-runtime	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	connect-test-plugins	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	connect-transforms	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	generator	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-clients	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.12	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-cloudera-metrics-reporter_2.13	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-cloudera-plugins	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-examples	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-group-coordinator	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-group-coordinator-api	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-log4j-appender	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-metadata	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-raft	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-server	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-server-common	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-shell	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-storage	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-storage-api	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-examples	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-scala_2.12	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-scala_2.13	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-test-utils	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0100	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0101	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0102	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-0110	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-10	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-11	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-20	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-21	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-22	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-23	3.9.1.7.3.2.0-957

Project	groupId	artifactId	version
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-24	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-25	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-26	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-27	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-28	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-30	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-31	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-32	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-33	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-34	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-35	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-36	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-37	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-streams-upgrade-system-tests-38	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-tools	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-tools-api	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka-transaction-coordinator	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka_2.12	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	kafka_2.13	3.9.1.7.3.2.0-957
Kafka	org.apache.kafka	trogdor	3.9.1.7.3.2.0-957
Knox	org.apache.knox	gateway-adapter	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-admin-ui	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-applications	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-cloud-bindings	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-demo-ldap	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-demo-ldap-launcher	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-discovery-ambari	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-discovery-cm	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-docker	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-i18n	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-i18n-logging-log4j	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-i18n-logging-sl4j	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-openapi-ui	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-performance-test	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-ha	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-identity-assertion-common	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-identity-assertion-concat	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-identity-assertion-hadoop-groups	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-identity-assertion-no-doas	2.1.0.7.3.2.0-957

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-provider-identity-assertion-pseudo	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-identity-assertion-regex	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-identity-assertion-switchcase	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-jersey	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-rewrite	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-rewrite-common	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-rewrite-func-hostmap-static	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-rewrite-func-inbound-query-param	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-rewrite-func-service-registry	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-rewrite-step-encrypt-uri	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-rewrite-step-secure-query	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-security-authc-anon	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-security-authc-remote	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-security-authz-acls	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-security-authz-composite	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-security-authz-path-acls	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-security-clientcert	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-security-hadoopauth	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-security-jwt	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-security-pac4j	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-security-preauth	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-security-shiro	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-provider-security-webappsec	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-release	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-server	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-server-launcher	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-server-xforwarded-filter	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-admin	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-as	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-auth	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-definitions	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-hashicorp-vault	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-hbase	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-health	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-hive	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-idbroker	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-idbroker-plugins	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-impala	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-jkg	2.1.0.7.3.2.0-957

Project	groupId	artifactId	version
Knox	org.apache.knox	gateway-service-knoxsso	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-knoxssout	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-knoxtoken	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-livy	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-metadata	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-nifi	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-nifi-registry	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-remoteconfig	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-rm	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-session	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-storm	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-test	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-tgs	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-vault	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-service-webhdfs	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-shell	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-shell-launcher	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-shell-release	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-shell-samples	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-spi	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-spi-common	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-test	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-test-idbroker	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-test-release-utils	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-test-utils	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-topology-hadoop-xml	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-topology-simple	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-util-common	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-util-configinjector	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-util-launcher	2.1.0.7.3.2.0-957
Knox	org.apache.knox	gateway-util-urltemplate	2.1.0.7.3.2.0-957
Knox	org.apache.knox	hadoop-examples	2.1.0.7.3.2.0-957
Knox	org.apache.knox	knox-cli-launcher	2.1.0.7.3.2.0-957
Knox	org.apache.knox	knox-homepage-ui	2.1.0.7.3.2.0-957
Knox	org.apache.knox	knox-token-generation-ui	2.1.0.7.3.2.0-957
Knox	org.apache.knox	knox-token-management-ui	2.1.0.7.3.2.0-957
Knox	org.apache.knox	knox-webshell-ui	2.1.0.7.3.2.0-957
Knox	org.apache.knox	webhdfs-kerb-test	2.1.0.7.3.2.0-957
Knox	org.apache.knox	webhdfs-test	2.1.0.7.3.2.0-957

Project	groupId	artifactId	version
Kudu	org.apache.kudu	kudu-backup-common	1.18.0.7.3.2.0-957
Kudu	org.apache.kudu	kudu-backup-tools	1.18.0.7.3.2.0-957
Kudu	org.apache.kudu	kudu-backup3_2.12	1.18.0.7.3.2.0-957
Kudu	org.apache.kudu	kudu-client	1.18.0.7.3.2.0-957
Kudu	org.apache.kudu	kudu-flatbuffers	1.18.0.7.3.2.0-957
Kudu	org.apache.kudu	kudu-hive	1.18.0.7.3.2.0-957
Kudu	org.apache.kudu	kudu-proto	1.18.0.7.3.2.0-957
Kudu	org.apache.kudu	kudu-replication	1.18.0.7.3.2.0-957
Kudu	org.apache.kudu	kudu-spark3-tools_2.12	1.18.0.7.3.2.0-957
Kudu	org.apache.kudu	kudu-spark3_2.12	1.18.0.7.3.2.0-957
Kudu	org.apache.kudu	kudu-subprocess	1.18.0.7.3.2.0-957
Kudu	org.apache.kudu	kudu-test-utils	1.18.0.7.3.2.0-957
Livy	org.apache.livy	livy-api	0.7.23000.7.3.2.0-957
Livy	org.apache.livy	livy-client-common	0.7.23000.7.3.2.0-957
Livy	org.apache.livy	livy-client-http	0.7.23000.7.3.2.0-957
Livy	org.apache.livy	livy-core_2.12	0.7.23000.7.3.2.0-957
Livy	org.apache.livy	livy-examples	0.7.23000.7.3.2.0-957
Livy	org.apache.livy	livy-integration-test	0.7.23000.7.3.2.0-957
Livy	org.apache.livy	livy-repl_2.12	0.7.23000.7.3.2.0-957
Livy	org.apache.livy	livy-rsc	0.7.23000.7.3.2.0-957
Livy	org.apache.livy	livy-scala-api_2.12	0.7.23000.7.3.2.0-957
Livy	org.apache.livy	livy-server	0.7.23000.7.3.2.0-957
Livy	org.apache.livy	livy-test-lib	0.7.23000.7.3.2.0-957
Livy	org.apache.livy	livy-thriftserver	0.7.23000.7.3.2.0-957
Livy	org.apache.livy	livy-thriftserver-session	0.7.23000.7.3.2.0-957
Lucene	org.apache.lucene	lucene-analyzers-common	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-analyzers-icu	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-analyzers-kuromoji	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-analyzers-morfologik	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-analyzers-nori	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-analyzers-opennlp	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-analyzers-phonetic	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-analyzers-smartcn	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-analyzers-stempel	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-backward-codecs	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-benchmark	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-classification	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-codecs	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-core	8.11.2.7.3.2.0-957

Project	groupId	artifactId	version
Lucene	org.apache.lucene	lucene-demo	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-expressions	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-facet	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-grouping	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-highlighter	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-join	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-memory	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-misc	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-monitor	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-queries	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-queryparser	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-replicator	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-sandbox	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-spatial-extras	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-spatial3d	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-suggest	8.11.2.7.3.2.0-957
Lucene	org.apache.lucene	lucene-test-framework	8.11.2.7.3.2.0-957
Oozie	org.apache.oozie	oozie-client	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-core	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-distro	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-examples	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-fluent-job-api	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-fluent-job-client	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-server	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-sharelib-distcp	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-sharelib-git	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-sharelib-hcatalog	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-sharelib-hive	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-sharelib-hive2	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-sharelib-oozie	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-sharelib-spark3	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-sharelib-sqoop	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-sharelib-streaming	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-tools	5.1.0.7.3.2.0-957
Oozie	org.apache.oozie	oozie-zookeeper-security-tests	5.1.0.7.3.2.0-957
ORC	org.apache.orc	orc-core	2.0.3.7.3.2.0-957
ORC	org.apache.orc	orc-mapreduce	2.0.3.7.3.2.0-957
ORC	org.apache.orc	orc-shims	2.0.3.7.3.2.0-957
ORC	org.apache.orc	orc-tools	2.0.3.7.3.2.0-957

Project	groupId	artifactId	version
Ozone	org.apache.ozone	hdds-annotation-processing	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-client	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-common	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-config	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-container-service	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-crypto-api	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-crypto-default	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-docs	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-erasurecode	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-interface-admin	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-interface-client	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-interface-server	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-managed-rocksdb	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-rocks-native	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-server-framework	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-server-scm	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	hdds-test-utils	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	mini-chaos-tests	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-cli-admin	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-cli-shell	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-client	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-common	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-csi	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-datanode	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-dev-support	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-dist	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-filesystem	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-filesystem-common	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-filesystem-hadoop2	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-filesystem-hadoop3	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-filesystem-shaded	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-freon	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-httpfsgateway	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-insight	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-integration-test	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-integration-test-recon	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-integration-test-s3	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-interface-client	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-interface-storage	2.0.0.7.3.2.0-957

Project	groupId	artifactId	version
Ozone	org.apache.ozone	ozone-manager	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-mini-cluster	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-network-tests	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-recon	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-reconcodegen	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-s3-secret-store	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-s3gateway	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	ozone-tools	2.0.0.7.3.2.0-957
Ozone	org.apache.ozone	rocksdb-checkpoint-differ	2.0.0.7.3.2.0-957
Parquet	org.apache.parquet	parquet-avro	1.15.2.7.3.2.0-957
Parquet	org.apache.parquet	parquet-column	1.15.2.7.3.2.0-957
Parquet	org.apache.parquet	parquet-common	1.15.2.7.3.2.0-957
Parquet	org.apache.parquet	parquet-encoding	1.15.2.7.3.2.0-957
Parquet	org.apache.parquet	parquet-format-structures	1.15.2.7.3.2.0-957
Parquet	org.apache.parquet	parquet-generator	1.15.2.7.3.2.0-957
Parquet	org.apache.parquet	parquet-hadoop	1.15.2.7.3.2.0-957
Parquet	org.apache.parquet	parquet-hadoop-bundle	1.15.2.7.3.2.0-957
Parquet	org.apache.parquet	parquet-jackson	1.15.2.7.3.2.0-957
Parquet	org.apache.parquet	parquet-pig	1.15.2.7.3.2.0-957
Parquet	org.apache.parquet	parquet-pig-bundle	1.15.2.7.3.2.0-957
Parquet	org.apache.parquet	parquet-protobuf	1.15.2.7.3.2.0-957
Parquet	org.apache.parquet	parquet-thrift	1.15.2.7.3.2.0-957
Parquet	org.apache.parquet	parquet-tools	1.15.2.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-client-embedded-hbase-2.6	5.2.1.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-client-lite-hbase-2.6	5.2.1.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-core	5.2.1.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-core-client	5.2.1.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-core-server	5.2.1.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.4.1	5.2.1.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.0	5.2.1.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.5.4	5.2.1.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-hbase-compat-2.6.0	5.2.1.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-mapreduce-byo-shaded-hbase-hbase-2.6	5.2.1.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-pherf	5.2.1.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-queryserver	6.0.0.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-queryserver-client	6.0.0.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-queryserver-it	6.0.0.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-queryserver-load-balancer	6.0.0.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-queryserver-orchestrator	6.0.0.7.3.2.0-957

Project	groupId	artifactId	version
Phoenix	org.apache.phoenix	phoenix-server-hbase-2.6	5.2.1.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix-tracing-webapp	5.2.1.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix5-hive	6.0.0.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix5-hive-shaded	6.0.0.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix5-spark3	6.0.0.7.3.2.0-957
Phoenix	org.apache.phoenix	phoenix5-spark3-shaded	6.0.0.7.3.2.0-957
Ranger	org.apache.ranger	conditions-enrichers	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	credentialbuilder	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	embeddedwebserver	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	hdfs-nn-perf	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	jisql	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ldapconfigcheck	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-adls-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-atlas-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-atlas-plugin-shim	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-authn	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-caii-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-clo-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-common-ha	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-distro	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-examples-distro	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-gs-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-hbase-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-hbase-plugin-shim	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-hdfs-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-hdfs-plugin-shim	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-hive-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-hive-plugin-shim	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-intg	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-kafka-connect-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-kafka-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-kafka-plugin-shim	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-kms	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-kms-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-kms-plugin-shim	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-knox-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-knox-plugin-shim	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-kudu-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-kylin-plugin	2.6.0.7.3.2.0-957

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-kylin-plugin-shim	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-metrics	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-nestedstructure-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-nifi-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-nifi-registry-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-opensearch-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-ozone-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-ozone-plugin-shim	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-plugin-classloader	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-plugin-perf	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-plugins-audit	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-plugins-common	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-plugins-cred	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-plugins-installer	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-policymigration	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-raz-adls	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-raz-chained-plugins	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-raz-hook-abfs	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-raz-hook-s3	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-raz-intg	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-raz-processor	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-raz-s3	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-raz-s3-lib	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-rms-common	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-rms-hive	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-rms-plugins-common	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-rms-tools	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-rms-webapp	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-s3-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-sampleapp-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-schema-registry-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-solr-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-solr-plugin-shim	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-sqoop-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-sqoop-plugin-shim	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-storm-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-storm-plugin-shim	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-tagsync	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-tools	2.6.0.7.3.2.0-957

Project	groupId	artifactId	version
Ranger	org.apache.ranger	ranger-trino-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-util	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-yarn-plugin	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ranger-yarn-plugin-shim	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	sample-client	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	sampleapp	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	ugsync-util	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	unixauthclient	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	unixauthservice	2.6.0.7.3.2.0-957
Ranger	org.apache.ranger	unixusersync	2.6.0.7.3.2.0-957
Solr	org.apache.solr	solr-analysis-extras	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-analytics	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-cell	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-core	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-dataimporthandler	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-dataimporthandler-extras	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-gcs-repository	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-jaegertracer-configurator	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-langid	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-ltr	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-prometheus-exporter	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-s3-repository	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-security-util	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-solrj	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-test-framework	8.11.2.7.3.2.0-957
Solr	org.apache.solr	solr-velocity	8.11.2.7.3.2.0-957
Spark	org.apache.spark	spark-avro_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-catalyst_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-common-utils_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-connect-client-jvm_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-connect-common_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-connect_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-core_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-graphx_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-hadoop-cloud_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-hive_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-kubernetes_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-kvstore_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-launcher_2.12	3.5.4.7.3.2.0-957

Project	groupId	artifactId	version
Spark	org.apache.spark	spark-mllib-local_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-mllib_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-network-common_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-network-shuffle_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-network-yarn_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-protobuf_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-repl_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-shaded-raz	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-sketch_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-sql-api_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-sql-kafka-0-10_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-sql_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-streaming-kafka-0-10-assembly_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-streaming-kafka-0-10_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-streaming_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-tags_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-token-provider-kafka-0-10_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-unsafe_2.12	3.5.4.7.3.2.0-957
Spark	org.apache.spark	spark-yarn_2.12	3.5.4.7.3.2.0-957
Sqoop	org.apache.sqoop	sqoop	1.4.7.7.3.2.0-957
Sqoop	org.apache.sqoop	sqoop-test	1.4.7.7.3.2.0-957
Tez	org.apache.tez	hadoop-shim	0.9.1.7.3.2.0-957
Tez	org.apache.tez	hadoop-shim-2.8	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-api	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-aux-services	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-build-tools	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-common	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-dag	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-examples	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-ext-service-tests	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-history-parser	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-javadoc-tools	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-job-analyzer	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-mapreduce	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-protobuf-history-plugin	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-runtime-internals	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-runtime-library	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-tests	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-yarn-timeline-cache-plugin	0.9.1.7.3.2.0-957

Project	groupId	artifactId	version
Tez	org.apache.tez	tez-yarn-timeline-history	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-yarn-timeline-history-with-acls	0.9.1.7.3.2.0-957
Tez	org.apache.tez	tez-yarn-timeline-history-with-fs	0.9.1.7.3.2.0-957
ZooKeeper	org.apache.zookeeper	zookeeper	3.8.5.7.3.2.0-957
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-fatjar	3.8.5.7.3.2.0-957
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-loggraph	3.8.5.7.3.2.0-957
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-rest	3.8.5.7.3.2.0-957
ZooKeeper	org.apache.zookeeper	zookeeper-contrib-zooinspector	3.8.5.7.3.2.0-957
ZooKeeper	org.apache.zookeeper	zookeeper-it	3.8.5.7.3.2.0-957
ZooKeeper	org.apache.zookeeper	zookeeper-jute	3.8.5.7.3.2.0-957
ZooKeeper	org.apache.zookeeper	zookeeper-prometheus-metrics	3.8.5.7.3.2.0-957
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-election	3.8.5.7.3.2.0-957
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-lock	3.8.5.7.3.2.0-957
ZooKeeper	org.apache.zookeeper	zookeeper-recipes-queue	3.8.5.7.3.2.0-957

Fixed Common Vulnerabilities and Exposures

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.2 release, and its Service Packs and Cumulative hotfixes.

Cloudera Runtime 7.3.2.20000 SP2

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.2.20000 SP2.

CVE ID	Description
CVE-2026-34040	Moby is an open source container framework. Prior to version 29.3.1, a security vulnerability has been detected that allows attackers to bypass authorization plugins (AuthZ). This issue has been patched in version 29.3.1.
CVE-2026-69248	Cryptography is a package designed to expose cryptographic primitives and recipes to Python developers. In versions 45.0.0 through 48.0.0, if an intermediate constrained CA permits the DNS name foo.example.com, and the leaf certificate has a wildcard in its DNS SAN of *.example.com, python-cryptography's verifier accepts which allows escaping outside of the permitted names. The core issue is in DNSConstraint::matches, where a wildcard pattern was treated as matching a more-specific permitted constraint even though *.example.com can expand to sibling names such as bar.example.com outside foo.example.com. This allows acceptance of an invalid certificate chain. This issue is fixed in 49.0.0.

Cloudera Runtime 7.3.2.10000 SP1

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.2.10000 SP1.

CVE ID	Description
CVE-2020-13956	Apache HttpClient versions prior to version 4.5.13 and 5.0.3 can misinterpret malformed authority component in request URIs passed to the library as java.net.URI object and pick the wrong target host for request execution.

CVE ID	Description
CVE-2023-0286	There is a type confusion vulnerability relating to X.400 address processing inside an X.509 GeneralName. X.400 addresses were parsed as an ASN1_STRING but the public structure definition for GENERAL_NAME incorrectly specified the type of the x400Address field as ASN1_TYPE. This field is subsequently interpreted by the OpenSSL function GENERAL_NAME_cmp as an ASN1_TYPE rather than an ASN1_STRING. When CRL checking is enabled (i.e. the application sets the X509_V_FLAG_CRL_CHECK flag), this vulnerability may allow an attacker to pass arbitrary pointers to a memcmp call, enabling them to read memory contents or enact a denial of service. In most cases, the attack requires the attacker to provide both the certificate chain and CRL, neither of which need to have a valid signature. If the attacker only controls one of these inputs, the other input must already contain an X.400 address as a CRL distribution point, which is uncommon. As such, this vulnerability is most likely to only affect applications which have implemented their own functionality for retrieving CRLs over a network.
CVE-2023-1428	There exists an vulnerability causing an abort() to be called in gRPC. The following headers cause gRPC's C++ implementation to abort() when called via http2: te: x (x != trailers) :scheme: x (x != http, https) grpclb_client_stats: x (x == anything) On top of sending one of those headers, a later header must be sent that gets the total header size past 8KB. We recommend upgrading past git commit 2485fa94bd8a723e5c977d55a3ce10b301b437f8 or v1.53 and above.
CVE-2023-23931	cryptography is a package designed to expose cryptographic primitives and recipes to Python developers. In affected versions `Cipher.update_into` would accept Python objects which implement the buffer protocol, but provide only immutable buffers. This would allow immutable objects (such as `bytes`) to be mutated, thus violating fundamental rules of Python and resulting in corrupted output. This now correctly raises an exception. This issue has been present since `update_into` was originally introduced in cryptography 1.8.
CVE-2023-49083	cryptography is a package designed to expose cryptographic primitives and recipes to Python developers. Calling `load_pem_pkcs7_certificates` or `load_der_pkcs7_certificates` could lead to a NULL-pointer dereference and segfault. Exploitation of this vulnerability poses a serious risk of Denial of Service (DoS) for any application attempting to deserialize a PKCS7 blob/certificate. The consequences extend to potential disruptions in system availability and stability. This vulnerability has been patched in version 41.0.6.
CVE-2023-50782	A flaw was found in the python-cryptography package. This issue may allow a remote attacker to decrypt captured messages in TLS servers that use RSA key exchanges, which may lead to exposure of confidential or sensitive data.
CVE-2024-0727	Issue summary: Processing a maliciously formatted PKCS12 file may lead OpenSSL to crash leading to a potential Denial of Service attack Impact summary: Applications loading files in the PKCS12 format from untrusted sources might terminate abruptly. A file in PKCS12 format can contain certificates and keys and may come from an untrusted source. The PKCS12 specification allows certain fields to be NULL, but OpenSSL does not correctly check for this case. This can lead to a NULL pointer dereference that results in OpenSSL crashing. If an application processes PKCS12 files from an untrusted source using the OpenSSL APIs then that application will be vulnerable to this issue. OpenSSL APIs that are vulnerable to this are: PKCS12_parse(), PKCS12_unpack_p7data(), PKCS12_unpack_p7encdata(), PKCS12_unpack_authsafes() and PKCS12_newpass(). We have also fixed a similar issue in SMIME_write_PKCS7(). However since this function is related to writing data we do not consider it security significant. The FIPS modules in 3.2, 3.1 and 3.0 are not affected by this issue.
CVE-2024-5569	A Denial of Service (DoS) vulnerability exists in the jaraco/zipp library, affecting all versions prior to 3.19.1. The vulnerability is triggered when processing a specially crafted zip file that leads to an infinite loop. This issue also impacts the zipfile module of CPython, as features from the third-party zipp library are later merged into CPython, and the affected code is identical in both projects. The infinite loop can be initiated through the use of functions affecting the `Path` module in both zipp and zipfile, such as `joinpath`, the overloaded division operator, and `iterdir`. Although the infinite loop is not resource exhaustive, it prevents the application from responding. The vulnerability was addressed in version 3.19.1 of jaraco/zipp.
CVE-2024-11392	Hugging Face Transformers MobileViTV2 Deserialization of Untrusted Data Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of Hugging Face Transformers. User interaction is required to exploit this vulnerability in that the target must visit a malicious page or open a malicious file. The specific flaw exists within the handling of configuration files. The issue results from the lack of proper validation of user-supplied data, which can result in deserialization of untrusted data. An attacker can leverage this vulnerability to execute code in the context of the current user. Was ZDI-CAN-24322.
CVE-2024-11393	Hugging Face Transformers MaskFormer Model Deserialization of Untrusted Data Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of Hugging Face Transformers. User interaction is required to exploit this vulnerability in that the target must visit a malicious page or open a malicious file. The specific flaw exists within the parsing of model files. The issue results from the lack of proper validation of user-supplied data, which can result in deserialization of untrusted data. An attacker can leverage this vulnerability to execute code in the context of the current user. Was ZDI-CAN-25191.

CVE ID	Description
CVE-2024-11394	Hugging Face Transformers Trax Model Deserialization of Untrusted Data Remote Code Execution Vulnerability. This vulnerability allows remote attackers to execute arbitrary code on affected installations of Hugging Face Transformers. User interaction is required to exploit this vulnerability in that the target must visit a malicious page or open a malicious file. The specific flaw exists within the handling of model files. The issue results from the lack of proper validation of user-supplied data, which can result in deserialization of untrusted data. An attacker can leverage this vulnerability to execute code in the context of the current user. Was ZDI-CAN-25012.
CVE-2024-12720	A Regular Expression Denial of Service (ReDoS) vulnerability was identified in the huggingface/transformers library, specifically in the file tokenization_nougat_fast.py. The vulnerability occurs in the post_process_single() function, where a regular expression processes specially crafted input. The issue stems from the regex exhibiting exponential time complexity under certain conditions, leading to excessive backtracking. This can result in significantly high CPU usage and potential application downtime, effectively creating a Denial of Service (DoS) scenario. The affected version is v4.46.3 (latest).
CVE-2024-21520	Versions of the package djangorestframework before 3.15.2 are vulnerable to Cross-site Scripting (XSS) via the break_long_headers template filter due to improper input sanitization before splitting and joining with tags.
CVE-2024-22513	djangorestframework-simplejwt version 5.3.1 and before is vulnerable to information disclosure. A user can access web application resources even after their account has been disabled due to missing user validation checks via the for_user method.
CVE-2024-26130	cryptography is a package designed to expose cryptographic primitives and recipes to Python developers. Starting in version 38.0.0 and prior to version 42.0.4, if `pkcs12.serialize_key_and_certificates` is called with both a certificate whose public key did not match the provided private key and an `encryption_algorithm` with `hmac_hash` set (via `PrivateFormat.PKCS12.encryption_builder().hmac_hash(...)`, then a NULL pointer dereference would occur, crashing the Python process. This has been resolved in version 42.0.4, the first version in which a `ValueError` is properly raised.
CVE-2024-45231	An issue was discovered in Django v5.1.1, v5.0.9, and v4.2.16. The django.contrib.auth.forms.PasswordResetForm class, when used in a view implementing password reset flows, allows remote attackers to enumerate user e-mail addresses by sending password reset requests and observing the outcome (only when e-mail sending is consistently failing).
CVE-2024-47081	Requests is a HTTP library. Due to a URL parsing issue, Requests releases prior to 2.32.4 may leak .netrc credentials to third parties for specific maliciously-crafted URLs. Users should upgrade to version 2.32.4 to receive a fix. For older versions of Requests, use of the .netrc file can be disabled with `trust_env=False` on one's Requests Session.
CVE-2024-47554	Uncontrolled Resource Consumption vulnerability in Apache Commons IO. The org.apache.commons.io.input.XmlStreamReader class may excessively consume CPU resources when processing maliciously crafted input. This issue affects Apache Commons IO: from 2.0 before 2.14.0. Users are recommended to upgrade to version 2.14.0 or later, which fixes the issue.
CVE-2025-1194	A Regular Expression Denial of Service (ReDoS) vulnerability was identified in the huggingface/transformers library, specifically in the file `tokenization_gpt_neox_japanese.py` of the GPT-NeoX-Japanese model. The vulnerability occurs in the SubWordJapaneseTokenizer class, where regular expressions process specially crafted inputs. The issue stems from a regex exhibiting exponential complexity under certain conditions, leading to excessive backtracking. This can result in high CPU usage and potential application downtime, effectively creating a Denial of Service (DoS) scenario. The affected version is v4.48.1 (latest).
CVE-2025-2099	A vulnerability in the `preprocess_string()` function of the `transformers.testing_utils` module in huggingface/transformers version v4.48.3 allows for a Regular Expression Denial of Service (ReDoS) attack. The regular expression used to process code blocks in docstrings contains nested quantifiers, leading to exponential backtracking when processing input with a large number of newline characters. An attacker can exploit this by providing a specially crafted payload, causing high CPU usage and potential application downtime, effectively resulting in a Denial of Service (DoS) scenario.
CVE-2025-2148	A vulnerability was found in PyTorch 2.6.0+cu124. It has been declared as critical. Affected by this vulnerability is the function torch.ops.profiler._call_end_callbacks_on_jit_fut of the component Tuple Handler. The manipulation of the argument None leads to memory corruption. The attack can be launched remotely. The complexity of an attack is rather high. The exploitation appears to be difficult.
CVE-2025-2149	A vulnerability was found in PyTorch 2.6.0+cu124. It has been rated as problematic. Affected by this issue is the function nnq_Sigmoid of the component Quantized Sigmoid Module. The manipulation of the argument scale/zero_point leads to improper initialization. The attack needs to be approached locally. The complexity of an attack is rather high. The exploitation is known to be difficult. The exploit has been disclosed to the public and may be used.

CVE ID	Description
CVE-2025-2953	A vulnerability, which was classified as problematic, has been found in PyTorch 2.6.0+cu124. Affected by this issue is the function torch.mkldnn_max_pool2d. The manipulation leads to denial of service. An attack has to be approached locally. The exploit has been disclosed to the public and may be used. The real existence of this vulnerability is still doubted at the moment. The security policy of the project warns to use unknown models which might establish malicious effects.
CVE-2025-2998	A vulnerability was found in PyTorch 2.6.0. It has been declared as critical. Affected by this vulnerability is the function torch.nn.utils.rnn.pad_packed_sequence. The manipulation leads to memory corruption. Local access is required to approach this attack. The exploit has been disclosed to the public and may be used.
CVE-2025-3263	A Regular Expression Denial of Service (ReDoS) vulnerability was discovered in the Hugging Face Transformers library, specifically in the `get_configuration_file()` function within the `transformers.configuration_utils` module. The affected version is 4.49.0, and the issue is resolved in version 4.51.0. The vulnerability arises from the use of a regular expression pattern `config\.(.*)\.json` that can be exploited to cause excessive CPU consumption through crafted input strings, leading to catastrophic backtracking. This can result in model serving disruption, resource exhaustion, and increased latency in applications using the library.
CVE-2025-3264	A Regular Expression Denial of Service (ReDoS) vulnerability was discovered in the Hugging Face Transformers library, specifically in the `get_imports()` function within `dynamic_module_utils.py`. This vulnerability affects versions 4.49.0 and is fixed in version 4.51.0. The issue arises from a regular expression pattern `\s*try\s*:.*?except.*?:` used to filter out try/except blocks from Python code, which can be exploited to cause excessive CPU consumption through crafted input strings due to catastrophic backtracking. This vulnerability can lead to remote code loading disruption, resource exhaustion in model serving, supply chain attack vectors, and development pipeline disruption.
CVE-2025-3730	A vulnerability, which was classified as problematic, was found in PyTorch 2.6.0. Affected is the function torch.nn.functional.ctc_loss of the file aten/src/ATen/native/LossCTC.cpp. The manipulation leads to denial of service. An attack has to be approached locally. The exploit has been disclosed to the public and may be used. The real existence of this vulnerability is still doubted at the moment. The name of the patch is 46fc5d8e360127361211cb237d5f9eef0223e567. It is recommended to apply a patch to fix this issue. The security policy of the project warns to use unknown models which might establish malicious effects.
CVE-2025-3777	Hugging Face Transformers versions up to 4.49.0 are affected by an improper input validation vulnerability in the `image_utils.py` file. The vulnerability arises from insecure URL validation using the `startswith()` method, which can be bypassed through URL username injection. This allows attackers to craft URLs that appear to be from YouTube but resolve to malicious domains, potentially leading to phishing attacks, malware distribution, or data exfiltration. The issue is fixed in version 4.52.1.
CVE-2025-3933	A Regular Expression Denial of Service (ReDoS) vulnerability was discovered in the Hugging Face Transformers library, specifically within the DonutProcessor class's `token2json()` method. This vulnerability affects versions 4.50.3 and earlier, and is fixed in version 4.52.1. The issue arises from the regex pattern `<s_(.*?)>` which can be exploited to cause excessive CPU consumption through crafted input strings due to catastrophic backtracking. This vulnerability can lead to service disruption, resource exhaustion, and potential API service vulnerabilities, impacting document processing tasks using the Donut model.
CVE-2025-5197	A Regular Expression Denial of Service (ReDoS) vulnerability exists in the Hugging Face Transformers library, specifically in the `convert_tf_weight_name_to_pt_weight_name()` function. This function, responsible for converting TensorFlow weight names to PyTorch format, uses a regex pattern `/[^/]*___([^/]*)/` that can be exploited to cause excessive CPU consumption through crafted input strings due to catastrophic backtracking. The vulnerability affects versions up to 4.51.3 and is fixed in version 4.53.0. This issue can lead to service disruption, resource exhaustion, and potential API service vulnerabilities, impacting model conversion processes between TensorFlow and PyTorch formats.
CVE-2025-6051	A Regular Expression Denial of Service (ReDoS) vulnerability was discovered in the Hugging Face Transformers library, specifically within the `normalize_numbers()` method of the `EnglishNormalizer` class. This vulnerability affects versions up to 4.52.4 and is fixed in version 4.53.0. The issue arises from the method's handling of numeric strings, which can be exploited using crafted input strings containing long sequences of digits, leading to excessive CPU consumption. This vulnerability impacts text-to-speech and number normalization tasks, potentially causing service disruption, resource exhaustion, and API vulnerabilities.
CVE-2025-6638	A Regular Expression Denial of Service (ReDoS) vulnerability was discovered in the Hugging Face Transformers library, specifically affecting the MarianTokenizer's `remove_language_code()` method. This vulnerability is present in version 4.52.4 and has been fixed in version 4.53.0. The issue arises from inefficient regex processing, which can be exploited by crafted input strings containing malformed language code patterns, leading to excessive CPU consumption and potential denial of service.

CVE ID	Description
CVE-2025-6921	The huggingface/transformers library, versions prior to 4.53.0, is vulnerable to Regular Expression Denial of Service (ReDoS) in the AdamWeightDecay optimizer. The vulnerability arises from the _do_use_weight_decay method, which processes user-controlled regular expressions in the include_in_weight_decay and exclude_from_weight_decay lists. Malicious regular expressions can cause catastrophic backtracking during the re.search call, leading to 100% CPU utilization and a denial of service. This issue can be exploited by attackers who can control the patterns in these lists, potentially causing the machine learning task to hang and rendering services unresponsive.
CVE-2025-11965	In Eclipse Vert.x versions [4.0.0, 4.5.21] and [5.0.0, 5.0.4], a StaticHandler configuration for restricting access to hidden files fails to restrict access to hidden directories, allowing unauthorized users to retrieve files within them (e.g. '.git/config').
CVE-2025-11966	In Eclipse Vert.x versions [4.0.0, 4.5.21] and [5.0.0, 5.0.4], when "directory listing" is enabled, file and directory names are inserted into generated HTML without proper escaping in the href, title, and link attributes. An attacker who can create or rename files or directories within a served path can craft filenames containing malicious script or HTML content, leading to stored cross-site scripting (XSS) that executes in the context of users viewing the affected directory listing.
CVE-2025-24970	Netty, an asynchronous, event-driven network application framework, has a vulnerability starting in version 4.1.91.Final and prior to version 4.1.118.Final. When a special crafted packet is received via SslHandler it doesn't correctly handle validation of such a packet in all cases which can lead to a native crash. Version 4.1.118.Final contains a patch. As workaround its possible to either disable the usage of the native SSLEngine or change the code manually.
CVE-2025-25193	Netty, an asynchronous, event-driven network application framework, has a vulnerability in versions up to and including 4.1.118.Final. An unsafe reading of environment file could potentially cause a denial of service in Netty. When loaded on an Windows application, Netty attempts to load a file that does not exist. If an attacker creates such a large file, the Netty application crash. A similar issue was previously reported as CVE-2024-47535. This issue was fixed, but the fix was incomplete in that null-bytes were not counted against the input limit. Commit d1fbda62d3a47835d3fb35db8bd42ecc205a5386 contains an updated fix.
CVE-2025-32434	PyTorch is a Python package that provides tensor computation with strong GPU acceleration and deep neural networks built on a tape-based autograd system. In version 2.5.1 and prior, a Remote Command Execution (RCE) vulnerability exists in PyTorch when loading a model using torch.load with weights_only=True. This issue has been patched in version 2.6.0.
CVE-2025-46551	JRuby-OpenSSL is an add-on gem for JRuby that emulates the Ruby OpenSSL native library. Starting in JRuby-OpenSSL version 0.12.1 and prior to version 0.15.4 (corresponding to JRuby versions starting in 9.3.4.0 prior to 9.4.12.1 and 10.0.0.0 prior to 10.0.0.1), when verifying SSL certificates, JRuby-OpenSSL does not verify that the hostname presented in the certificate matches the one the user tries to connect to. This means a man-in-the-middle could just present any valid cert for a completely different domain they own, and JRuby would accept the cert. Anybody using JRuby to make requests of external APIs, or scraping the web, that depends on https to connect securely. JRuby-OpenSSL version 0.15.4 contains a fix for the issue. This fix is included in JRuby versions 10.0.0.1 and 9.4.12.1.
CVE-2025-47287	Tornado is a Python web framework and asynchronous networking library. When Tornado's ``multipart/form-data`` parser encounters certain errors, it logs a warning but continues trying to parse the remainder of the data. This allows remote attackers to generate an extremely high volume of logs, constituting a DoS attack. This DoS is compounded by the fact that the logging subsystem is synchronous. All versions of Tornado prior to 6.5.0 are affected. The vulnerable parser is enabled by default. Upgrade to Tornado version 6.50 to receive a patch. As a workaround, risk can be mitigated by blocking `Content-Type: multipart/form-data` in a proxy.
CVE-2025-48432	An issue was discovered in Django 5.2 before 5.2.3, 5.1 before 5.1.11, and 4.2 before 4.2.23. Internal HTTP response logging does not escape request.path, which allows remote attackers to potentially manipulate log output via crafted URLs. This may lead to log injection or forgery when logs are viewed in terminals or processed by external systems.

CVE ID	Description
CVE-2025-48734	Improper Access Control vulnerability in Apache Commons. A special BeanIntrospector class was added in version 1.9.2. This can be used to stop attackers from using the declared class property of Java enum objects to get access to the classloader. However this protection was not enabled by default. PropertyUtilsBean (and consequently BeanUtilsBean) now disallows declared class level property access by default. Releases 1.11.0 and 2.0.0-M2 address a potential security issue when accessing enum properties in an uncontrolled way. If an application using Commons BeanUtils passes property paths from an external source directly to the getProperty() method of PropertyUtilsBean, an attacker can access the enum's class loader via the "declaredClass" property available on all Java "enum" objects. Accessing the enum's "declaredClass" allows remote attackers to access the ClassLoader and execute arbitrary code. The same issue exists with PropertyUtilsBean.getNestedProperty(). Starting in versions 1.11.0 and 2.0.0-M2 a special BeanIntrospector suppresses the "declaredClass" property. Note that this new BeanIntrospector is enabled by default, but you can disable it to regain the old behavior; see section 2.5 of the user's guide and the unit tests. This issue affects Apache Commons BeanUtils 1.x before 1.11.0, and 2.x before 2.0.0-M2.Users of the artifact commons-beanutils:commons-beanutils 1.x are recommended to upgrade to version 1.11.0, which fixes the issue. Users of the artifact org.apache.commons:commons-beanutils2 2.x are recommended to upgrade to version 2.0.0-M2, which fixes the issue.
CVE-2025-49128	Jackson-core contains core low-level incremental ("streaming") parser and generator abstractions used by Jackson Data Processor. Starting in version 2.0.0 and prior to version 2.13.0, a flaw in jackson-core's `JsonLocation._appendSourceDesc` method allows up to 500 bytes of unintended memory content to be included in exception messages. When parsing JSON from a byte array with an offset and length, the exception message incorrectly reads from the beginning of the array instead of the logical payload start. This results in possible information disclosure in systems using pooled or reused buffers, like Netty or Vert.x. This issue was silently fixed in jackson-core version 2.13.0, released on September 30, 2021, via PR #652. All users should upgrade to version 2.13.0 or later. If upgrading is not immediately possible, applications can mitigate the issue by disabling exception message exposure to clients to avoid returning parsing exception messages in HTTP responses and/or disabling source inclusion in exceptions to prevent Jackson from embedding any source content in exception messages, avoiding leakage.
CVE-2025-54920	This issue affects Apache Spark: before 3.5.7 and 4.0.1. Users are recommended to upgrade to version 3.5.7 or 4.0.1 and above, which fixes the issue. Summary Apache Spark 3.5.4 and earlier versions contain a code execution vulnerability in the Spark History Web UI due to overly permissive Jackson deserialization of event log data. This allows an attacker with access to the Spark event logs directory to inject malicious JSON payloads that trigger deserialization of arbitrary classes, enabling command execution on the host running the Spark History Server. Details The vulnerability arises because the Spark History Server uses Jackson polymorphic deserialization with @JsonTypeInfo.Id.CLASS on SparkListenerEvent objects, allowing an attacker to specify arbitrary class names in the event JSON. This behavior permits instantiating unintended classes, such as org.apache.hive.jdbc.HiveConnection, which can perform network calls or other malicious actions during deserialization. The attacker can exploit this by injecting crafted JSON content into the Spark event log files, which the History Server then deserializes on startup or when loading event logs. For example, the attacker can force the History Server to open a JDBC connection to a remote attacker-controlled server, demonstrating remote command injection capability. Proof of Concept: 1. Run Spark with event logging enabled, writing to a writable directory (spark-logs). 2. Inject the following JSON at the beginning of an event log file: { "Event": "org.apache.hive.jdbc.HiveConnection", "uri": "jdbc:hive2://<IP>:<PORT>/", "info": { "hive.metastore.uris": "thrift://<IP>:<PORT>" } } 3. Start the Spark History Server with logs pointing to the modified directory. 4. The Spark History Server initiates a JDBC connection to the attacker's server, confirming the injection. Impact An attacker with write access to Spark event logs can execute arbitrary code on the server running the History Server, potentially compromising the entire system.
CVE-2025-57833	An issue was discovered in Django 4.2 before 4.2.24, 5.1 before 5.1.12, and 5.2 before 5.2.6. FilteredRelation is subject to SQL injection in column aliases, using a suitably crafted dictionary, with dictionary expansion, as the **kwargs passed QuerySet.annotate() or QuerySet.alias().
CVE-2025-58056	Netty is an asynchronous event-driven network application framework for development of maintainable high performance protocol servers and clients. In versions 4.1.124.Final, and 4.2.0.Alpha3 through 4.2.4.Final, Netty incorrectly accepts standalone newline characters (LF) as a chunk-size line terminator, regardless of a preceding carriage return (CR), instead of requiring CRLF per HTTP/1.1 standards. When combined with reverse proxies that parse LF differently (treating it as part of the chunk extension), attackers can craft requests that the proxy sees as one request but Netty processes as two, enabling request smuggling attacks. This is fixed in versions 4.1.125.Final and 4.2.5.Final.
CVE-2025-59059	Remote Code Execution Vulnerability in NashornScriptEngineCreator is reported in Apache Ranger versions <= 2.7.0. Users are recommended to upgrade to version 2.8.0, which fixes this issue.
CVE-2025-59060	Hostname verification bypass issue in Apache Ranger NiFiRegistryClient/NiFiClient is reported in Apache Ranger versions <= 2.7.0. Users are recommended to upgrade to version 2.8.0, which fixes this issue.

CVE ID	Description
CVE-2025-64458	An issue was discovered in 5.1 before 5.1.14, 4.2 before 4.2.26, and 5.2 before 5.2.8. NFKC normalization in Python is slow on Windows. As a consequence, `django.http.HttpResponseRedirect`, `django.http.HttpResponsePermanentRedirect`, and the shortcut `django.shortcuts.redirect` were subject to a potential denial-of-service attack via certain inputs with a very large number of Unicode characters. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank Seokchan Yoon for reporting this issue.
CVE-2025-64459	An issue was discovered in 5.1 before 5.1.14, 4.2 before 4.2.26, and 5.2 before 5.2.8. The methods `QuerySet.filter()`, `QuerySet.exclude()`, and `QuerySet.get()`, and the class `Q()`, are subject to SQL injection when using a suitably crafted dictionary, with dictionary expansion, as the `_connector` argument. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank cyberstan for reporting this issue.
CVE-2025-67724	Tornado is a Python web framework and asynchronous networking library. In versions 6.5.2 and below, the supplied reason phrase is used unescaped in HTTP headers (where it could be used for header injection) or in HTML in the default error page (where it could be used for XSS) and can be exploited by passing untrusted or malicious data into the reason argument. Used by both RequestHandler.set_status and tornado.web.HTTPError, the argument is designed to allow applications to pass custom "reason" phrases (the "Not Found" in HTTP/1.1 404 Not Found) to the HTTP status line (mainly for non-standard status codes). This issue is fixed in version 6.5.3.
CVE-2025-67725	Tornado is a Python web framework and asynchronous networking library. In versions 6.5.2 and below, a single maliciously crafted HTTP request can block the server's event loop for an extended period, caused by the HTTPHeaders.add method. The function accumulates values using string concatenation when the same header name is repeated, causing a Denial of Service (DoS). Due to Python string immutability, each concatenation copies the entire string, resulting in $O(n^2)$ time complexity. The severity can vary from high if max_header_size has been increased from its default, to low if it has its default value of 64KB. This issue is fixed in version 6.5.3.
CVE-2025-67726	Tornado is a Python web framework and asynchronous networking library. Versions 6.5.2 and below use an inefficient algorithm when parsing parameters for HTTP header values, potentially causing a DoS. The _parseparam function in httputil.py is used to parse specific HTTP header values, such as those in multipart/form-data and repeatedly calls string.count() within a nested loop while processing quoted semicolons. If an attacker sends a request with a large number of maliciously crafted parameters in a Content-Disposition header, the server's CPU usage increases quadratically ($O(n^2)$) during parsing. Due to Tornado's single event loop architecture, a single malicious request can cause the entire server to become unresponsive for an extended period. This issue is fixed in version 6.5.3.
CVE-2025-68161	The Socket Appender in Apache Log4j Core versions 2.0-beta9 through 2.25.2 does not perform TLS hostname verification of the peer certificate, even when the verifyHostName https://logging.apache.org/log4j/2.x/manual/appenders/network.html#SslConfiguration-attr-verifyHostName configuration attribute or the log4j2.sslVerifyHostName https://logging.apache.org/log4j/2.x/manual/systemproperties.html#log4j2.sslVerifyHostName system property is set to true. This issue may allow a man-in-the-middle attacker to intercept or redirect log traffic under the following conditions: * The attacker is able to intercept or redirect network traffic between the client and the log receiver. * The attacker can present a server certificate issued by a certification authority trusted by the Socket Appender's configured trust store (or by the default Java trust store if no custom trust store is configured). Users are advised to upgrade to Apache Log4j Core version 2.25.3, which addresses this issue. As an alternative mitigation, the Socket Appender may be configured to use a private or restricted trust root to limit the set of trusted certificates.
CVE-2026-3505	Allocation of resources without limits or throttling, Uncontrolled Resource Consumption vulnerability in Legion of the Bouncy Castle Inc. BC-JAVA bcpg on all (pg modules). This vulnerability is associated with program files AEADEncDataPacket.Java, BcAEADUtil.Java, JceAEADUtil.Java, OperatorHelper.Java. This issue affects BC-JAVA: from 1.74 before 1.80.2, from 1.81 before 1.81.1, from 1.82 before 1.84.
CVE-2026-5598	Covert timing channel vulnerability in Legion of the Bouncy Castle Inc. BC-JAVA core on all (core modules). This vulnerability is associated with program files FrodoEngine.Java. This issue affects BC-JAVA: from 1.71 before 1.80.2, from 1.81 before 1.81.1, from 1.82 before 1.84.
CVE-2026-24880	Inconsistent Interpretation of HTTP Requests ('HTTP Request/Response Smuggling') vulnerability in Apache Tomcat via invalid chunk extension. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.18, from 10.1.0-M1 through 10.1.52, from 9.0.0.M1 through 9.0.115, from 8.5.0 through 8.5.100, from 7.0.0 through 7.0.109. Other, unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.20, 10.1.52 or 9.0.116, which fix the issue.
CVE-2026-25854	Occasional URL redirection to untrusted Site ('Open Redirect') vulnerability in Apache Tomcat via the LoadBalancerDrainingValve. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.18, from 10.1.0-M1 through 10.1.52, from 9.0.0.M23 through 9.0.115, from 8.5.30 through 8.5.100. Other, unsupported versions may also be affected Users are recommended to upgrade to version 11.0.20, 10.1.53 or 9.0.116, which fix the issue.
CVE-2026-25990	Pillow is a Python imaging library. From 10.3.0 to before 12.1.1, an out-of-bounds write may be triggered when loading a specially crafted PSD image. This vulnerability is fixed in 12.1.1.

CVE ID	Description
CVE-2026-27727	mchange-commons-java, a library that provides Java utilities, includes code that mirrors early implementations of JNDI functionality, including support for remote `factoryClassLocation` values, by which code can be downloaded and invoked within a running application. If an attacker can provoke an application to read a maliciously crafted `jaxax.naming.Reference` or serialized object, they can provoke the download and execution of malicious code. Implementations of this functionality within the JDK were disabled by default behind a System property that defaults to `false`, `com.sun.jndi.ldap.object.trustURLCodebase`. However, since mchange-commons-java includes an independent implementation of JNDI derefencing, libraries (such as c3p0) that resolve references via that implementation could be provoked to download and execute malicious code even after the JDK was hardened. Mirroring the JDK patch, mchange-commons-java's JNDI functionality is gated by configuration parameters that default to restrictive values starting in version 0.4.0. No known workarounds are available. Versions prior to 0.4.0 should be avoided on application CLASSPATHs.
CVE-2026-29129	Configured cipher preference order not preserved vulnerability in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.16 through 11.0.18, from 10.1.51 through 10.1.52, from 9.0.114 through 9.0.115. Users are recommended to upgrade to version 11.0.20, 10.1.53 or 9.0.116, which fix the issue.
CVE-2026-29145	CLIENT_CERT authentication does not fail as expected for some scenarios when soft fail is disabled vulnerability in Apache Tomcat, Apache Tomcat Native. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.18, from 10.1.0-M7 through 10.1.52, from 9.0.83 through 9.0.115; Apache Tomcat Native: from 1.1.23 through 1.1.34, from 1.2.0 through 1.2.39, from 1.3.0 through 1.3.6, from 2.0.0 through 2.0.13. Users are recommended to upgrade to version Tomcat Native 1.3.7 or 2.0.14 and Tomcat 11.0.20, 10.1.53 and 9.0.116, which fix the issue.
CVE-2026-31958	Tornado is a Python web framework and asynchronous networking library. In versions of Tornado prior to 6.5.5, the only limit on the number of parts in multipart/form-data is the max_body_size setting (default 100MB). Since parsing occurs synchronously on the main thread, this creates the possibility of denial-of-service due to the cost of parsing very large multipart bodies with many parts. This vulnerability is fixed in 6.5.5.
CVE-2026-32597	PyJWT is a JSON Web Token implementation in Python. Prior to 2.12.0, PyJWT does not validate the crit (Critical) Header Parameter defined in RFC 7515 §4.1.11. When a JWS token contains a crit array listing extensions that PyJWT does not understand, the library accepts the token instead of rejecting it. This violates the MUST requirement in the RFC. This vulnerability is fixed in 2.12.0.
CVE-2026-32990	Improper Input Validation vulnerability in Apache Tomcat due to an incomplete fix of CVE-2025-66614. This issue affects Apache Tomcat: from 11.0.15 through 11.0.19, from 10.1.50 through 10.1.52, from 9.0.113 through 9.0.115. Users are recommended to upgrade to version 11.0.20, 10.1.53 or 9.0.116, which fix the issue.
CVE-2026-33701	OpenTelemetry Java Instrumentation provides OpenTelemetry auto-instrumentation and instrumentation libraries for Java. In versions prior to 2.26.1, the RMI instrumentation registered a custom endpoint that deserialized incoming data without applying serialization filters. On JDK version 16 and earlier, an attacker with network access to a JMX or RMI port on an instrumented JVM could exploit this to potentially achieve remote code execution. All three of the following conditions must be true to exploit this vulnerability: First, OpenTelemetry Java instrumentation is attached as a Java agent (`-javaagent`) on Java 16 or earlier. Second, JMX/RMI port has been explicitly configured via `-Dcom.sun.management.jmxremote.port` and is network-reachable. Third, gadget-chain-compatible library is present on the classpath. This results in arbitrary remote code execution with the privileges of the user running the instrumented JVM. For JDK >= 17, no action is required, but upgrading is strongly encouraged. For JDK < 17, upgrade to version 2.26.1 or later. As a workaround, set the system property `-Dotel.instrumentation.rmi.enabled=false` to disable the RMI integration.
CVE-2026-34477	The fix for CVE-2025-68161 https://logging.apache.org/security.html#CVE-2025-68161 was incomplete: it addressed hostname verification only when enabled via the log4j2.sslVerifyHostName https://logging.apache.org/log4j/2.x/manual/systemproperties.html#log4j2.sslVerifyHostName system property, but not when configured through the verifyHostName https://logging.apache.org/log4j/2.x/manual/appenders/network.html#SslConfiguration-attr-verifyHostName attribute of the <Ssl> element. Although the verifyHostName configuration attribute was introduced in Log4j Core 2.12.0, it was silently ignored in all versions through 2.25.3, leaving TLS connections vulnerable to interception regardless of the configured value. A network-based attacker may be able to perform a man-in-the-middle attack when all of the following conditions are met: * An SMTP, Socket, or Syslog appender is in use. * TLS is configured via a nested <Ssl> element. * The attacker can present a certificate issued by a CA trusted by the appender's configured trust store, or by the default Java trust store if none is configured. This issue does not affect users of the HTTP appender, which uses a separate verifyHostname https://logging.apache.org/log4j/2.x/manual/appenders/network.html#HttpAppender-attr-verifyHostName attribute that was not subject to this bug and verifies host names by default. Users are advised to upgrade to Apache Log4j Core 2.25.4, which corrects this issue.

CVE ID	Description
CVE-2026-34480	Apache Log4j Core's XmlLayout https://logging.apache.org/log4j/2.x/manual/layouts.html#XmlLayout , in versions up to and including 2.25.3, fails to sanitize characters forbidden by the XML 1.0 specification https://www.w3.org/TR/xml/#charsets producing invalid XML output whenever a log message or MDC value contains such characters. The impact depends on the StAX implementation in use: * JRE built-in StAX: Forbidden characters are silently written to the output, producing malformed XML. Conforming parsers must reject such documents with a fatal error, which may cause downstream log-processing systems to drop the affected records. * Alternative StAX implementations (e.g., Woodstox https://github.com/FasterXML/woodstox , a transitive dependency of the Jackson XML Dataformat module): An exception is thrown during the logging call, and the log event is never delivered to its intended appender, only to Log4j's internal status logger. Users are advised to upgrade to Apache Log4j Core 2.25.4, which corrects this issue by sanitizing forbidden characters before XML output.
CVE-2026-34481	Apache Log4j's JsonTemplateLayout https://logging.apache.org/log4j/2.x/manual/json-template-layout.html , in versions up to and including 2.25.3, produces invalid JSON output when log events contain non-finite floating-point values (NaN, Infinity, or -Infinity), which are prohibited by RFC 8259. This may cause downstream log processing systems to reject or fail to index affected records. An attacker can exploit this issue only if both of the following conditions are met: * The application uses JsonTemplateLayout. * The application logs a MapMessage containing an attacker-controlled floating-point value. Users are advised to upgrade to Apache Log4j JSON Template Layout 2.25.4, which corrects this issue.
CVE-2026-34483	Improper Encoding or Escaping of Output vulnerability in the JsonAccessLogValve component of Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.20, from 10.1.0-M1 through 10.1.53, from 9.0.40 through 9.0.116. Users are recommended to upgrade to version 11.0.21, 10.1.54 or 9.0.117 , which fix the issue.
CVE-2026-34487	Insertion of Sensitive Information into Log File vulnerability in the cloud membership for clustering component of Apache Tomcat exposed the Kubernetes bearer token. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.20, from 10.1.0-M1 through 10.1.53, from 9.0.13 through 9.0.116. Users are recommended to upgrade to version 11.0.21, 10.1.54 or 9.0.117, which fix the issue.
CVE-2026-34500	CLIENT_CERT authentication does not fail as expected for some scenarios when soft fail is disabled and FFM is used in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M14 through 11.0.20, from 10.1.22 through 10.1.53, from 9.0.92 through 9.0.116. Users are recommended to upgrade to version 11.0.21, 10.1.54 or 9.0.117, which fixes the issue.
CVE-2026-35536	In Tornado before 6.5.5, cookie attribute injection could occur because the domain, path, and samesite arguments to .RequestHandler.set_cookie were not checked for crafted characters.
CVE-2026-41284	Allocation of Resources Without Limits or Throttling vulnerability in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.21, from 10.1.0-M1 through 10.1.54, from 9.0.0.M1 through 9.0.117. Older, unsupported versions may also be affected. Users are recommended to upgrade to version [FIXED_VERSION], which fixes the issue.
CVE-2026-41293	Improper Input Validation vulnerability in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.21, from 10.1.0-M1 through 10.1.54, from 9.0.0.M1 through 9.0.117, from 10.0.0-M1 through 10.0.27. Older, end of support versions may also be affected. Users are recommended to upgrade to version [FIXED_VERSION], which fixes the issue.
CVE-2026-41409	The fix for CVE-2024-52046 in Apache MINA AbstractIoBuffer.getObject() was incomplete. The classname allowlist of classes allowed to be deserialized was applied too late after a static initializer in a class to be read might already have been executed. Affected versions are Apache MINA 2.0.0 <= 2.0.27, 2.1.0 <= 2.1.10, and 2.2.0 <= 2.2.5. The problem is resolved in Apache MINA 2.0.28, 2.1.11, and 2.2.6 by applying the classname allowlist earlier. Affected are applications using Apache MINA that call IoBuffer.getObject(). Applications using Apache MINA are advised to upgrade
CVE-2026-41635	Apache MINA's AbstractIoBuffer.resolveClass() contains two branches, one of them (for static classes or primitive types) does not check the class at all, bypassing the classname allowlist and allowing arbitrary code to be executed. The fix checks if the class is present in the accepted class filter before calling Class.forName(). Affected versions are Apache MINA 2.0.0 <= 2.0.27, 2.1.0 <= 2.1.10, and 2.2.0 <= 2.2.5. The problem is resolved in Apache MINA 2.0.28, 2.1.11, and 2.2.6 by applying the classname allowlist earlier. Affected are applications using Apache MINA that call IoBuffer.getObject(). Applications using Apache MINA are advised to upgrade.
CVE-2026-42304	Twisted is an event-based framework for internet applications, supporting Python 3.6+. Prior to 26.4.0rc2, the twisted.names module is vulnerable to a Denial of Service (DoS) attack via resource exhaustion during DNS name decompression. A remote, unauthenticated attacker can exploit this by sending a crafted TCP DNS packet containing deeply chained compression pointers. This flaw bypasses previous loop-prevention logic, causing the single-threaded Twisted reactor to hang while processing millions of recursive lookups, effectively freezing the server. This vulnerability is fixed in 26.4.0rc2.

CVE ID	Description
CVE-2026-42498	Exposure of HTTP Authentication Header to unexpected hosts during WebSocket authentication vulnerability in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.21, from 10.1.0-M1 through 10.1.54, from 9.0.2 through 9.0.117, from 8.5.24 through 8.5.100, from 7.0.83 through 7.0.109. Users are recommended to upgrade to version 11.0.22, 10.1.55 or 9.0.118, which fix the issue.
CVE-2026-42577	Netty is an asynchronous, event-driven network application framework. From 4.2.0.Final to 4.2.13.Final , Netty's epoll transport fails to detect and close TCP connections that receive a RST after being half-closed, leading to stale channels that are never cleaned up and, in some code paths, a 100% CPU busy-loop in the event loop thread. This vulnerability is fixed in 4.2.13.Final.
CVE-2026-42578	Netty is an asynchronous, event-driven network application framework. Prior to 4.2.13.Final and 4.1.133.Final, Netty's HttpProxyHandler constructs HTTP CONNECT requests with header validation explicitly disabled. The newInitialMessage() method creates headers using DefaultHttpHeadersFactory.headersFactory().withValidation(false), then adds user-provided outboundHeaders without any CRLF validation. This allows an attacker who can influence the outbound headers to inject arbitrary HTTP headers into the CONNECT request sent to the proxy server. This vulnerability is fixed in 4.2.13.Final and 4.1.133.Final.
CVE-2026-42586	Netty is an asynchronous, event-driven network application framework. Prior to 4.2.13.Final and 4.1.133.Final, the Netty Redis codec encoder (RedisEncoder) writes user-controlled string content directly to the network output buffer without validating or sanitizing CRLF (\r\n) characters. Since the Redis Serialization Protocol (RESP) uses CRLF as the command/response delimiter, an attacker who can control the content of a Redis message can inject arbitrary Redis commands or forge fake responses. This vulnerability is fixed in 4.2.13.Final and 4.1.133.Final.
CVE-2026-42778	The fix for CVE-2026-41409 was not applied to the 2.1.X and 2.2.X branches. Here was the original issue description: The fix for CVE-2024-52046 in Apache MINA AbstractIoBuffer.getObject() was incomplete. The classname allowlist of classes allowed to be deserialized was applied too late after a static initializer in a class to be read might already have been executed. Affected versions are Apache MINA 2.1.0 <= 2.1.11, and 2.2.0 <= 2.2.6. The problem is resolved in Apache MINA 2.1.12, and 2.2.7 by applying the classname allowlist earlier. Affected are applications using Apache MINA that call IoBuffer.getObject(). Applications using Apache MINA are advised to upgrade The fix for CVE-2024-52046 in Apache MINA AbstractIoBuffer.getObject() was incomplete. The classname allowlist of classes allowed to be deserialized was applied too late after a static initializer in a class to be read might already have been executed. Affected versions are Apache MINA 2.1.0 <= 2.1.110, and 2.2.0 <= 2.2.6. The problem is resolved in Apache MINA 2.1.12, and 2.2.7 by applying the classname allowlist earlier. Affected are applications using Apache MINA that call IoBuffer.getObject(). Applications using Apache MINA are advised to upgrade
CVE-2026-42779	The fix for CVE-2026-41635 was not applied to the 2.1.X and 2.2.X branches. Here was the original issue description: Apache MINA's AbstractIoBuffer.resolveClass() contains two branches, one of them (for static classes or primitive types) does not check the class at all, bypassing the classname allowlist and allowing arbitrary code to be executed. The fix checks if the class is present in the accepted class filter before calling Class.forName(). Affected versions are Apache MINA 2.1.0 <= 2.1.11, and 2.2.0 <= 2.2.6. The problem is resolved in Apache MINA 2.1.12, and 2.2.7 by applying the classname allowlist earlier. Affected are applications using Apache MINA that call IoBuffer.getObject(). Applications using Apache MINA are advised to upgrade.
CVE-2026-43512	DEPRECATED: Authentication Bypass Issues vulnerability in digest authentication in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.21, from 10.1.0-M1 through 10.1.54, from 9.0.0.M1 through 9.0.117, from 8.5.0 through 8.5.100, from before 7.0.0. Older unsupported versions any also be affect Users are recommended to upgrade to version 11.0.22, 10.1.55 or 9.0.118 which fix the issue.
CVE-2026-43513	Improper Handling of Case Sensitivity vulnerability in LockOutRealm in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.21, from 10.1.0-M1 through 10.1.54, from 9.0.0.M1 through 9.0.117, from 8.5.0 through 8.5.100, from 7.0.0 through 7.0.109. Older unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.22, 10.1.55 or 9.0.118 which fix the issue.
CVE-2026-43514	Observable Timing Discrepancy vulnerability when comparing AJP secret in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.21, from 10.1.0-M1 through 10.1.54, from 9.0.0.M1 through 9.0.117, from 8.5.0 through 8.5.100, from 7.0.0 through 7.0.109. Older unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.22, 10.1.55 or 9.0.118 which fix the issue.
CVE-2026-43515	Improper Authorization vulnerability when multiple method constraints define an HTTP method for the same extension in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.21, from 10.1.0-M1 through 10.1.54, from 9.0.0.M1 through 9.0.117, from 8.5.0 through 8.5.100, from 7.0.0 through 7.0.109. Users are recommended to upgrade to version 11.0.22, 10.1.55 or 9.0.118 which fix the issue.

CVE ID	Description
CVE-2026-44248	Netty is an asynchronous, event-driven network application framework. Prior to 4.2.13.Final and 4.1.133.Final, the MQTT 5 header Properties section is parsed and buffered before any message size limit is applied. Specifically, in MqttDecoder, the decodeVariableHeader() method is called before the bytesRemainingBeforeVariableHeader > maxBytesInMessage check. The decodeVariableHeader() can call other methods which will call decodeProperties(). Effectively, Netty does not apply any limits to the size of the properties being decoded. Additionally, because MqttDecoder extends ReplayingDecoder, Netty will repeatedly re-parse the enormous Properties sections and buffer the bytes in memory, until the entire thing parses to completion. This can cause high resource usage in both CPU and memory. This vulnerability is fixed in 4.2.13.Final and 4.1.133.Final.
CVE-2026-44250	Netty is a network application framework for development of protocol servers and clients. In netty-codec-redis prior to versions 4.1.135.Final and 4.2.15.Final, an attacker can cause DoS by sending a crafted Redis payload with deeply nested arrays. This forces the server to allocate a massive number of state objects and collections, leading to memory exhaustion and an OutOfMemoryError. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-44545	daphne before 4.2.2 did not pass maxFramePayloadSize or maxMessagePayloadSize to Autobahn's WebSocketServerFactory. Because Autobahn defaults both values to 0 (unlimited), an unauthenticated remote attacker could send arbitrarily large WebSocket messages or frames, causing excessive memory consumption and a denial of service.
CVE-2026-44546	daphne before 4.2.2 reconstructs a raw HTTP request from Twisted's parsed headers and feeds it to autobahn for WebSocket handshake processing. Twisted does not treat \x0b, \x0c, \x1c, \x1d, \x1e, or \x85 as header line separators, but autobahn decodes header values to str and calls splitlines(). An attacker can exploit this parser differential to inject additional headers into the ASGI scope passed to the application. daphne now rejects requests with these bytes in any header value with a 400 response.
CVE-2026-44660	UltraJSON is a fast JSON encoder and decoder written in pure C with bindings for Python 3.7+. Prior to 5.12.1, when ujson.dump() writes to a file-like object and the write operation raises an exception, the serialized JSON string object is not decremented, leaking memory. Each failed write operation leaks the full size of the serialized payload. This vulnerability is fixed in 5.12.1.
CVE-2026-44890	Netty is a network application framework for development of protocol servers and clients. In netty-codec-redis prior to versions 4.1.135.Final and 4.2.15.Final, an attacker can cause DoS by sending crafted Redis payloads across multiple connections without `\r\n`. This exhausts the server's direct memory pool (OutOfDirectMemoryError), preventing legitimate connections from being processed. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-44891	Netty is a network application framework for development of protocol servers and clients. Prior to 4.1.136.Final and 4.2.16.Final, io.netty.handler.codec.stomp.StompSubframeDecoder fails to limit the total number of headers or their cumulative size per frame, and the maxLineLength parameter only restricts individual header lines. An attacker can send a large number of short headers that are accumulated in memory inside DefaultStompHeadersSubframe until the JVM throws an OutOfMemoryError, causing denial of service for servers exposing a STOMP endpoint based on StompSubframeDecoder. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-44893	Netty is a network application framework for development of protocol servers and clients. In netty-codec-haproxy prior to versions 4.1.135.Final and 4.2.15.Final, when decoding a PP2_TYPE_SSL TLV, HAProxyMessage.readNextTLV() first calls `header.retainedSlice(header.readerIndex(), length)` and only then reads the 1-byte client field and 4-byte verify field. If the attacker sets the TLV length below 5, the subsequent readByte/readInt throws IndexOutOfBoundsException. HAProxyMessageDecoder only catches HAProxyProtocolException around this call, so the IOOBE propagates and the retained slice on the pooled cumulation buffer is never released. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-46340	Netty is a network application framework for development of protocol servers and clients. In versions of netty-transport-sctp prior to 4.1.135.Final and 4.2.15.Final, for each non-complete SctpMessage fragment the handler does `fragments.put(streamId, Unpooled.wrappedBuffer(frag, byteBuf))`, wrapping the previous accumulator and the new slice into a *new* CompositeByteBuf every time. After N fragments the accumulator is an N-deep chain of composites, each holding references and component arrays; readableBytes()/getBytes() on the final buffer recurse N levels. There is no limit on N, on total bytes, or on the number of streamIdentifiers an attacker can open (each gets its own map entry). A peer that never sets the `complete` flag can grow this structure indefinitely from tiny 1-byte DATA chunks. Versions 4.1.135.Final and 4.2.15.Final patch the issue.

CVE ID	Description
CVE-2026-47065	ZDRES-232: resolveProxyClass Not Overridden - acceptMatchers Filter Bypass via java.lang.reflect.Proxy Assessment: Fully addressed. When the serialised stream contains a TC_PROXYCLASSDESC (the marker for a java.lang.reflect.Proxy), JDK's ObjectInputStream.readProxyDesc() is dispatched. JDK then calls the default ObjectInputStream.resolveProxyClass(interfaces) implementation, which performs Class.forName(intf, false, latestUserDefinedLoader()) for EACH interface name and constructs the proxy class â€" bypassing the accepted classes list . ZDRES-233: Class.forName(name, initialize=true, classLoader) in readClassDescriptor Triggers Static Initialiser of Allow-Listed Classes Assessment: Fully addressed. For ANY class on the allow-list, deserialising a stream that names it triggers the class's (static initialiser) BEFORE any instance is constructed. This means an attacker who supplies a class name on the allow-list (e.g., the developer wrote accept("com.myapp.*") , attacker supplies com.myapp.SomeClass) causes <clinit> of SomeClass â€" and many real-world classes have side-effecting static initialisers Both issues have been fixed.
CVE-2026-48006	Netty is a network application framework for development of protocol servers and clients. Prior to versions 4.1.135.Final and 4.2.15.Final, the RedisArrayAggregator handler permanently leaks pooled direct-memory buffers when a Redis pipeline connection closes before a RESP array aggregate completes. The handler retains child messages in per-handler state (`depths` field) but defines no `channelInactive`, `handlerRemoved`, or `exceptionCaught` method to release them when the pipeline tears down. Because the leaked buffers are slices of `PooledByteBufAllocator` chunks, they prevent those chunks from being returned to the JVM-wide direct-memory pool. Repeated connection churn by any network peer monotonically drains this shared pool, eventually causing allocation failures on all Netty channels in the process. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-48059	Netty is a network application framework for development of protocol servers and clients. Prior to versions 4.1.135.Final and 4.2.15.Final, the HAProxy PROXY protocol v2 codec in netty leaks native or heap memory on every connection when a client sends a syntactically valid header containing nested `PP2_TYPE_SSL` TLVs (type-length-value records) at depth two or greater. The leak occurs on the successful parse path — no exception is thrown, the message fires downstream, the decoder removes itself, and the application releases the `HAProxyMessage` normally. Yet the underlying cumulation buffer (a pooled, potentially direct `ByteBuf` allocated by the channel) remains permanently pinned. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-48522	PyJWT is a JSON Web Token implementation in Python. Prior to 2.13.0, PyJWKClient passes its uri argument directly to urllib.request.urlopen() which uses Python stdlib's default OpenerDirector registering HTTPHandler, HTTPSHandler, FTPHandler, FileHandler, and DataHandler. There is currently no documented option to restrict which schemes PyJWKClient will fetch. If an application's jku URL ingestion path accepts attacker-influenced URLs (e.g., from JWT header, configuration file, OAuth flow parameter), the attacker can cause PyJWKClient to read arbitrary local files via file:// (SSRF on local filesystem), cause PyJWKClient to attempt FTP / data-URI fetches (broader SSRF surface), or forge tokens that PyJWT verifies as valid. The library does not directly return non-HTTP(S) URI contents to the attacker; the chained "plant a JWKS to forge tokens" scenario described in the original report requires additional application-layer flaws (attacker write access to a filesystem path, untrusted jku derivation) that this fix does not address. This vulnerability is fixed in 2.13.0.
CVE-2026-48524	PyJWT is a JSON Web Token implementation in Python. Prior to 2.13.0, PyJWKClient.get_signing_key() forces a fresh HTTP request to the JWKS endpoint for every JWT with an unknown kid value, with no rate limiting. Since kid comes from the unverified token header, an attacker can trigger unlimited outbound requests. The vulnerability surfaces only when a JWKS fetch fails; an attacker can attempt to provoke that with sustained unknown-kid traffic, but the outcome depends on upstream JWKS-endpoint behavior (rate limiting, transient errors) which is beyond the attacker's control. This vulnerability is fixed in 2.13.0.
CVE-2026-48526	PyJWT is a JSON Web Token implementation in Python. Prior to 2.13.0, when the verifier is decoding JSON Web Tokens, while supporting both asymmetric and HMAC algorithms, the library does not validate use of JSON Web Keys in HMAC algorithm, allowing attacker to use the issuer public key as the secret key for HMAC algorithm. This vulnerability is fixed in 2.13.0.
CVE-2026-49853	Tornado is a Python web framework and asynchronous networking library. Prior to 6.5.6, SimpleAsyncHTTPClient shallow-copied redirected requests and removed only the Host header, leaving Authorization, auth_username, auth_password, and auth_mode in place when a redirect changed scheme, host, or port. This issue is fixed in version 6.5.6.
CVE-2026-49854	Tornado is a Python web framework and asynchronous networking library. Prior to 6.5.6, the optional native extension tornado.speedups implemented websocket_mask without validating that the mask argument is exactly four bytes, allowing the C function to read up to three bytes beyond the provided buffer when reached through Tornado XSRF token decoding with the native extension active. This issue is fixed in version 6.5.6.
CVE-2026-49855	Tornado is a Python web framework and asynchronous networking library. Prior to 6.5.6, Tornado gzip decompression routines processed limited-size chunks but did not enforce an overall limit on accumulated decompressed chunks, allowing a malicious server accessed by SimpleAsyncHTTPClient or an HTTPServer configured with decompress_request=True to consume effectively unlimited memory. This issue is fixed in version 6.5.6.

CVE ID	Description
CVE-2026-50011	Netty is a network application framework for development of protocol servers and clients. Prior to versions 4.1.135.Final and 4.2.15.Final, RedisArrayAggregator pre-allocates ArrayList with initial capacity equal to the RESP array element count declared in an array header. That count is taken from the wire before the corresponding child messages exist. A small malicious header can claim a huge initial capacity. Versions 4.1.135.Final and 4.2.15.Final patch the issue.
CVE-2026-50193	jackson-databind contains the general-purpose data-binding functionality and tree-model for Jackson Data Processor. From 2.13.0 until 2.14.0, a potential Denial-of-Service exists when attacker sends deeply nested JSON if (and only if) the service reads deeply nested (1000s of levels) JSON as JsonNode (ObjectMapper.readTree()) and writes out same (or modifided) node using JsonNode.toString(). This can consume significant amount of resources with concurrent relatively small requests (1000 nested arrays is 2kB). This vulnerability is fixed in 2.14.0.
CVE-2026-54911	No description available.
CVE-2026-55760	No description available.
CVE-2026-55851	Netty is a network application framework for development of protocol servers and clients. In versions 4.2.0.Final up to (but not including) 4.2.16.Final, and 4.1.0.Final up to (but not including) 4.1.135, the `HAProxyMessageDecoder` in Netty's `codec-haproxy` module performs protocol version detection by reading the 13th byte as a signed Java `byte` and widening it to `int` without masking; a PROXY protocol v2 binary prefix followed by version byte `0xFF` sign-extends to `-1`, collides with the decoder's need-more-data sentinel, and causes `ByteToMessageDecoder` to accumulate inbound bytes in an unbounded `cumulation` buffer until direct memory is exhausted. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-56817	Netty is a network application framework for development of protocol servers and clients. In versions 4.2.0.Final through 4.2.15.Final and 4.1.0.Final through 4.1.135.Final, any caller that can deliver bytes to a Netty channel pipeline containing `XmlDecoder` can send XML with a `DOCTYPE` declaration to an `AsyncXMLInputFactory` instantiated with no security configuration, leaving DTD and entity handling active depending on Aalto XML async parser behavior and creating conditional XML external entity risk. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-56818	Netty is an asynchronous, event-driven network application framework. Prior to 4.1.136.Final and 4.2.16.Final, the RedisArrayAggregator Redis codec clears retained partial aggregate state when the maxNestedArrayDepth limit is exceeded, but it does not clear the same state when the sibling maxElements limit is exceeded. A peer can start a valid RESP array, send a bulk string child, then send a nested array header longer than the configured maxElements. Netty throws a decoder exception in decodeRedisArrayHeader, but the existing partial aggregate remains retained in the handler. If the application leaves the channel alive after the exception, later messages are still consumed into the pre-error aggregate, allowing an unauthenticated peer to keep attacker-controlled aggregate state alive across a security-limit exception and pin retained pooled buffers. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-56820	Netty is a network application framework for development of protocol servers and clients. In versions 4.2.0.Final through 4.2.15.Final and prior to 4.1.135.Final, `OcspClient` does not validate that the `CertificateID` in an OCSP response matches the requested `CertificateID`, which can lead to replay attack. `OcspClient.validateResponse` accepts a legitimately signed `GOOD` status response for an unrelated certificate issued by the same CA, allowing bypass of revocation checks for another certificate. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-56821	Netty is an asynchronous, event-driven network application framework. Prior to versions 4.1.136.Final and 4.2.16.Final, the OcspServerCertificateValidator flags an out-of-date OCSP response but does not stop processing it, so an expired GOOD response is still reported as VALID, letting an on-path attacker replay a stale GOOD response to bypass revocation of a since-revoked certificate. Exploitation can lead to certificate revocation bypass via replay of an expired OCSP response. Any application using OcspServerCertificateValidator is affected; a revoked certificate can be accepted. This issue has been fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-56822	Netty is an asynchronous, event-driven network application framework. Prior to versions 4.1.136.Final and 4.2.16.Final, the OcspServerCertificateValidator forwards the SslHandshakeCompletionEvent before the asynchronous OCSP validation completes. This allows the client's downstream handlers to send sensitive application data (e.g., HTTP requests) to a revoked server before the channel is closed by the OCSP check. n io.netty.handler.ssl.ocsp.OcspServerCertificateValidator#userEventTriggered, when an SslHandshakeCompletionEvent is received, the validator immediately calls ctx.fireUserEventTriggered(evt). It then initiates an asynchronous OCSP query using OcspClient.query. Because the handshake completion event is forwarded immediately, downstream handlers in the client's pipeline are notified that the TLS handshake is successful. They may then begin reading and processing incoming application data or sending outgoing data. If the OCSP response later indicates the server's certificate is REVOKED, the validator closes the channel, but by this time, the client may have already leaked sensitive data to a revoked server or processed malicious responses from it. This issue has been fixed in versions 4.1.136.Final and 4.2.16.Final.

CVE ID	Description
CVE-2026-59884	pyasn1 is a generic ASN.1 library for Python. Prior to 0.6.4, the BER decoder shared by the CER and DER codecs parses long-form tags by accumulating continuation octets without an upper bound on the tag ID size, allowing a crafted input to force construction of an arbitrarily large integer with CPU cost growing quadratically and to trigger unhandled ValueError exceptions in Python 3.11+ error formatting paths. Any application decoding untrusted BER, CER, or DER input is affected. This issue is fixed in version 0.6.4.
CVE-2026-59885	pyasn1 is a generic ASN.1 library for Python. Prior to 0.6.4, the BER, CER, and DER decoders process OBJECT IDENTIFIER and RELATIVE-OID values in quadratic time relative to the number of arcs, so a small crafted payload containing an OID with many arcs consumes excessive CPU per decode() call and can deny service to applications that decode untrusted ASN.1 data. The corresponding encoders have the same quadratic behavior when an application re-encodes previously decoded attacker-supplied values. This issue is fixed in version 0.6.4.
CVE-2026-59886	pyasn1 is a generic ASN.1 library for Python. Prior to 0.6.4, the univ.Real type converted its mantissa, base, and exponent value to a Python float using exact big-integer exponentiation. A BER, CER, or DER encoded REAL value only a few bytes long can carry a very large exponent, causing float conversion through prettyPrint(), str(), comparison, arithmetic, int(), or an explicit float() call to consume excessive CPU and memory and hang applications that decode untrusted ASN.1 data and then print, log, or compare decoded objects. This issue is fixed in version 0.6.4.
CVE-2026-59919	Netty is an asynchronous, event-driven network application framework. In versions prior to 4.1.136.Final and 4.2.16.Final, Netty's HAProxy encoder (HAProxyMessageEncoder) writes AF_UNIX source and destination socket addresses into the HAProxy V1 text protocol without validating them for CRLF characters, so an attacker who controls an AF_UNIX address can inject \r\n sequences and split the single PROXY header into multiple lines. This is possible because the V1 protocol uses CRLF as its line terminator and, unlike IPv4/IPv6 addresses whose format checks implicitly reject CRLF, AF_UNIX addresses are only validated for length (up to 108 bytes), allowing a forged second PROXY header line that spoofs the client source/destination IP to a downstream server or load balancer. The issue is fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-59920	Netty is an asynchronous, event-driven network application framework. In versions prior to 4.1.136.Final and 4.2.16.Final, Netty's STOMP encoder (StompSubframeEncoder) does not escape or validate header values in CONNECT and CONNECTED frames, so raw newline (\n) characters in a header value are written directly to the wire, allowing an attacker who controls a header value to inject additional STOMP headers. This happens because the encoder intentionally skips escaping for CONNECT/CONNECTED frames per the STOMP 1.2 specification but never rejects the raw newlines, and since a broker parses each line as a separate header, an attacker controlling a value such as a user-supplied login or passcode can overwrite connection parameters or add authentication/role headers to bypass authentication or escalate privileges (the actual impact is broker-dependent). The issue is fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2026-59939	httplib2 is a comprehensive HTTP client library for Python. Prior to 0.32.0, httplib2 performs unbounded decompression of HTTP response bodies encoded with Content-Encoding: gzip or deflate in _decompressContent in httplib2/init.py, allowing a malicious or compromised HTTP server to return a small compressed payload that expands to an arbitrarily large size in memory and causes MemoryError or OOM-kill in the client process. This issue is fixed in version 0.32.0.
CVE-2026-73507	Netty is an asynchronous, event-driven network application framework. Prior to 4.1.136.Final and 4.2.16.Final, io.netty.handler.codec.xml.XmlFrameDecoder.decode() failed to preserve closing-tag parser state across invocations, so an unauthenticated remote attacker could trickle-feed repeated </ sequences that repeatedly rescanned the accumulated buffer and exhausted an EventLoop thread's CPU, causing denial of service with a maxFrameLength of 1 MB. This issue is fixed in versions 4.1.136.Final and 4.2.16.Final.
CVE-2025-12183	Out-of-bounds memory operations in org.lz4:lz4-java 1.8.0 and earlier allow remote attackers to cause denial of service and read adjacent memory via untrusted compressed input. Spark 3 is updated to address this issue.
CVE-2025-54920	Apache Spark 3.5.4 and earlier can allow code execution in the Spark History Web UI when deserializing crafted event log data. Spark 3 is updated to Apache Spark 3.5.8, which addresses this issue.
CVE-2025-33042	Improper Control of Generation of Code ('Code Injection') vulnerability in Apache Avro Java SDK when generating specific records from untrusted Avro schemas. This issue affects Apache Avro Java SDK: all versions through 1.11.4 and version 1.12.0. Users are recommended to upgrade to version 1.12.1 or 1.11.5, which fix the issue.
CVE-2026-42579	Netty is an asynchronous, event-driven network application framework. Prior to 4.2.13.Final and 4.1.133.Final, Netty's DNS codec does not enforce RFC 1035 domain name constraints during either encoding or decoding. This creates a bidirectional attack surface: malicious DNS responses can exploit the decoder, and user-influenced hostnames can exploit the encoder. This vulnerability is fixed in 4.2.13.Final and 4.1.133.Final.
CVE-2026-43869	Improper Validation of Certificate with Host Mismatch vulnerability in Apache Thrift. This issue affects Apache Thrift: before 0.23.0. Users are recommended to upgrade to version 0.23.0, which fixes the issue.

Cloudera Runtime 7.3.2.100 CHF 1

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.2.100 CHF 1.

Table 15: CVE IDs and description

CVE ID	Description
CVE-2026-24308	Improper handling of configuration values in ZKConfig in Apache ZooKeeper 3.8.5 and 3.9.4 on all platforms allows an attacker to expose sensitive information stored in client configuration in the client's logfile. Configuration values are exposed at INFO level logging rendering potential production systems affected by the issue. Users are recommended to upgrade to version 3.8.6 or 3.9.5 which fixes this issue.
CVE-2026-24281	Hostname verification in Apache ZooKeeper ZKTrustManager falls back to reverse DNS (PTR) when IP SAN validation fails, allowing attackers who control or spoof PTR records to impersonate ZooKeeper servers or clients with a valid certificate for the PTR name. It's important to note that attacker must present a certificate which is trusted by ZKTrustManager which makes the attack vector harder to exploit. Users are recommended to upgrade to version 3.8.6 or 3.9.5, which fixes this issue by introducing a new configuration option to disable reverse DNS lookup in client and quorum protocols.
CVE-2026-34480	Apache Log4j Core's XmlLayout https://logging.apache.org/log4j/2.x/manual/layouts.html#XmlLayout , in versions up to and including 2.25.3, fails to sanitize characters forbidden by the XML 1.0 specification https://www.w3.org/TR/xml/#charsets producing invalid XML output whenever a log message or MDC value contains such characters. The impact depends on the StAX implementation in use: * JRE built-in StAX: Forbidden characters are silently written to the output, producing malformed XML. Conforming parsers must reject such documents with a fatal error, which may cause downstream log-processing systems to drop the affected records. * Alternative StAX implementations (e.g., Woodstox https://github.com/FasterXML/woodstox , a transitive dependency of the Jackson XML Dataformat module): An exception is thrown during the logging call, and the log event is never delivered to its intended appender, only to Log4j's internal status logger. Users are advised to upgrade to Apache Log4j Core 2.25.4, which corrects this issue by sanitizing forbidden characters before XML output.
CVE-2025-67721	Aircompressor is a library with ports of the Snappy, LZO, LZ4, and Zstandard compression algorithms to Java. In versions 3.3 and below, incorrect handling of malformed data in Java-based decompressor implementations for Snappy and LZ4 allow remote attackers to read previous buffer contents via crafted compressed input. With certain crafted compressed inputs, elements from the output buffer can end up in the uncompressed output, potentially leaking sensitive data. This is relevant for applications that reuse the same output buffer to uncompress multiple inputs. This can be the case of a web server that allocates a fix-sized buffer for performance purposes. There is similar vulnerability in GHSA-cmp6-m4wj-q63q. This issue is fixed in version 3.4.
CVE-2023-0286	There is a type confusion vulnerability relating to X.400 address processing inside an X.509 GeneralName. X.400 addresses were parsed as an ASN1_STRING but the public structure definition for GENERAL_NAME incorrectly specified the type of the x400Address field as ASN1_TYPE. This field is subsequently interpreted by the OpenSSL function GENERAL_NAME_cmp as an ASN1_TYPE rather than an ASN1_STRING. When CRL checking is enabled (i.e. the application sets the X509_V_FLAG_CRL_CHECK flag), this vulnerability may allow an attacker to pass arbitrary pointers to a memcmp call, enabling them to read memory contents or enact a denial of service. In most cases, the attack requires the attacker to provide both the certificate chain and CRL, neither of which need to have a valid signature. If the attacker only controls one of these inputs, the other input must already contain an X.400 address as a CRL distribution point, which is uncommon. As such, this vulnerability is most likely to only affect applications which have implemented their own functionality for retrieving CRLs over a network.
CVE-2023-23931	cryptography is a package designed to expose cryptographic primitives and recipes to Python developers. In affected versions `Cipher.update_into` would accept Python objects which implement the buffer protocol, but provide only immutable buffers. This would allow immutable objects (such as `bytes`) to be mutated, thus violating fundamental rules of Python and resulting in corrupted output. This now correctly raises an exception. This issue has been present since `update_into` was originally introduced in cryptography 1.8.
CVE-2023-49083	cryptography is a package designed to expose cryptographic primitives and recipes to Python developers. Calling `load_pem_pkcs7_certificates` or `load_der_pkcs7_certificates` could lead to a NULL-pointer dereference and segfault. Exploitation of this vulnerability poses a serious risk of Denial of Service (DoS) for any application attempting to deserialize a PKCS7 blob/certificate. The consequences extend to potential disruptions in system availability and stability. This vulnerability has been patched in version 41.0.6.
CVE-2023-50782	A flaw was found in the python-cryptography package. This issue may allow a remote attacker to decrypt captured messages in TLS servers that use RSA key exchanges, which may lead to exposure of confidential or sensitive data.

CVE ID	Description
CVE-2024-0727	Issue summary: Processing a maliciously formatted PKCS12 file may lead OpenSSL to crash leading to a potential Denial of Service attack Impact summary: Applications loading files in the PKCS12 format from untrusted sources might terminate abruptly. A file in PKCS12 format can contain certificates and keys and may come from an untrusted source. The PKCS12 specification allows certain fields to be NULL, but OpenSSL does not correctly check for this case. This can lead to a NULL pointer dereference that results in OpenSSL crashing. If an application processes PKCS12 files from an untrusted source using the OpenSSL APIs then that application will be vulnerable to this issue. OpenSSL APIs that are vulnerable to this are: PKCS12_parse(), PKCS12_unpack_p7data(), PKCS12_unpack_p7encdata(), PKCS12_unpack_authsafes() and PKCS12_newpass(). We have also fixed a similar issue in SMIME_write_PKCS7(). However since this function is related to writing data we do not consider it security significant. The FIPS modules in 3.2, 3.1 and 3.0 are not affected by this issue.
CVE-2024-21520	Versions of the package djangorestframework before 3.15.2 are vulnerable to Cross-site Scripting (XSS) via the break_long_headers template filter due to improper input sanitization before splitting and joining with tags.
CVE-2024-22513	djangorestframework-simplejwt version 5.3.1 and before is vulnerable to information disclosure. A user can access web application resources even after their account has been disabled due to missing user validation checks via the for_user method.
CVE-2024-26130	cryptography is a package designed to expose cryptographic primitives and recipes to Python developers. Starting in version 38.0.0 and prior to version 42.0.4, if `pkcs12.serialize_key_and_certificates` is called with both a certificate whose public key did not match the provided private key and an `encryption_algorithm` with `hmac_hash` set (via `PrivateFormat.PKCS12.encryption_builder().hmac_hash(...)`, then a NULL pointer dereference would occur, crashing the Python process. This has been resolved in version 42.0.4, the first version in which a `ValueError` is properly raised.
CVE-2024-45231	An issue was discovered in Django v5.1.1, v5.0.9, and v4.2.16. The django.contrib.auth.forms.PasswordResetForm class, when used in a view implementing password reset flows, allows remote attackers to enumerate user e-mail addresses by sending password reset requests and observing the outcome (only when e-mail sending is consistently failing).
CVE-2024-47554	Uncontrolled Resource Consumption vulnerability in Apache Commons IO. The org.apache.commons.io.input.XmlStreamReader class may excessively consume CPU resources when processing maliciously crafted input. This issue affects Apache Commons IO: from 2.0 before 2.14.0. Users are recommended to upgrade to version 2.14.0 or later, which fixes the issue.
CVE-2025-47287	Tornado is a Python web framework and asynchronous networking library. When Tornado's ``multipart/form-data`` parser encounters certain errors, it logs a warning but continues trying to parse the remainder of the data. This allows remote attackers to generate an extremely high volume of logs, constituting a DoS attack. This DoS is compounded by the fact that the logging subsystem is synchronous. All versions of Tornado prior to 6.5.0 are affected. The vulnerable parser is enabled by default. Upgrade to Tornado version 6.50 to receive a patch. As a workaround, risk can be mitigated by blocking `Content-Type: multipart/form-data` in a proxy.
CVE-2025-48432	An issue was discovered in Django 5.2 before 5.2.3, 5.1 before 5.1.11, and 4.2 before 4.2.23. Internal HTTP response logging does not escape request.path, which allows remote attackers to potentially manipulate log output via crafted URLs. This may lead to log injection or forgery when logs are viewed in terminals or processed by external systems.
CVE-2025-49128	Jackson-core contains core low-level incremental ("streaming") parser and generator abstractions used by Jackson Data Processor. Starting in version 2.0.0 and prior to version 2.13.0, a flaw in jackson-core's `JsonLocation._appendSourceDesc` method allows up to 500 bytes of unintended memory content to be included in exception messages. When parsing JSON from a byte array with an offset and length, the exception message incorrectly reads from the beginning of the array instead of the logical payload start. This results in possible information disclosure in systems using pooled or reused buffers, like Netty or Vert.x. This issue was silently fixed in jackson-core version 2.13.0, released on September 30, 2021, via PR #652. All users should upgrade to version 2.13.0 or later. If upgrading is not immediately possible, applications can mitigate the issue by disabling exception message exposure to clients to avoid returning parsing exception messages in HTTP responses and/or disabling source inclusion in exceptions to prevent Jackson from embedding any source content in exception messages, avoiding leakage.
CVE-2025-57833	An issue was discovered in Django 4.2 before 4.2.24, 5.1 before 5.1.12, and 5.2 before 5.2.6. FilteredRelation is subject to SQL injection in column aliases, using a suitably crafted dictionary, with dictionary expansion, as the **kwargs passed QuerySet.annotate() or QuerySet.alias().
CVE-2025-64458	An issue was discovered in 5.1 before 5.1.14, 4.2 before 4.2.26, and 5.2 before 5.2.8. NFKC normalization in Python is slow on Windows. As a consequence, `django.http.HttpResponseRedirect`, `django.http.HttpResponsePermanentRedirect`, and the shortcut `django.shortcuts.redirect` were subject to a potential denial-of-service attack via certain inputs with a very large number of Unicode characters. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank Seokchan Yoon for reporting this issue.

CVE ID	Description
CVE-2025-64459	An issue was discovered in 5.1 before 5.1.14, 4.2 before 4.2.26, and 5.2 before 5.2.8. The methods `QuerySet.filter()`, `QuerySet.exclude()`, and `QuerySet.get()`, and the class `Q()`, are subject to SQL injection when using a suitably crafted dictionary, with dictionary expansion, as the `_connector` argument. Earlier, unsupported Django series (such as 5.0.x, 4.1.x, and 3.2.x) were not evaluated and may also be affected. Django would like to thank cyberstan for reporting this issue.
CVE-2026-4539	A security flaw has been discovered in pygments up to 2.19.2. The impacted element is the function AdlLexer of the file pygments/lexers/archetype.py. The manipulation results in inefficient regular expression complexity. The attack is only possible with local access. The exploit has been released to the public and may be used for attacks. The project was informed of the problem early through an issue report but has not responded yet.
CVE-2026-5598	Covert timing channel vulnerability in Legion of the Bouncy Castle Inc. BC-JAVA core on all (core modules). This vulnerability is associated with program files FrodoEngine.Java. This issue affects BC-JAVA: from 1.71 before 1.80.2, from 1.81 before 1.81.1, from 1.82 before 1.84.
CVE-2026-24880	Inconsistent Interpretation of HTTP Requests ('HTTP Request/Response Smuggling') vulnerability in Apache Tomcat via invalid chunk extension. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.18, from 10.1.0-M1 through 10.1.52, from 9.0.0.M1 through 9.0.115, from 8.5.0 through 8.5.100, from 7.0.0 through 7.0.109. Other, unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.20, 10.1.52 or 9.0.116, which fix the issue.
CVE-2026-25854	Occasional URL redirection to untrusted Site ('Open Redirect') vulnerability in Apache Tomcat via the LoadBalancerDrainingValve. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.18, from 10.1.0-M1 through 10.1.52, from 9.0.0.M23 through 9.0.115, from 8.5.30 through 8.5.100. Other, unsupported versions may also be affected Users are recommended to upgrade to version 11.0.20, 10.1.53 or 9.0.116, which fix the issue.
CVE-2026-27727	mchange-commons-java, a library that provides Java utilities, includes code that mirrors early implementations of JNDI functionality, including support for remote `factoryClassLocation` values, by which code can be downloaded and invoked within a running application. If an attacker can provoke an application to read a maliciously crafted `jaxax.naming.Reference` or serialized object, they can provoke the download and execution of malicious code. Implementations of this functionality within the JDK were disabled by default behind a System property that defaults to `false`, `com.sun.jndi.ldap.object.trustURLCodebase`. However, since mchange-commons-java includes an independent implementation of JNDI derefencing, libraries (such as c3p0) that resolve references via that implementation could be provoked to download and execute malicious code even after the JDK was hardened. Mirroring the JDK patch, mchange-commons-java's JNDI functionality is gated by configuration parameters that default to restrictive values starting in version 0.4.0. No known workarounds are available. Versions prior to 0.4.0 should be avoided on application CLASSPATHs.
CVE-2026-29129	Configured cipher preference order not preserved vulnerability in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.16 through 11.0.18, from 10.1.51 through 10.1.52, from 9.0.114 through 9.0.115. Users are recommended to upgrade to version 11.0.20, 10.1.53 or 9.0.116, which fix the issue.
CVE-2026-29145	CLIENT_CERT authentication does not fail as expected for some scenarios when soft fail is disabled vulnerability in Apache Tomcat, Apache Tomcat Native. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.18, from 10.1.0-M7 through 10.1.52, from 9.0.83 through 9.0.115; Apache Tomcat Native: from 1.1.23 through 1.1.34, from 1.2.0 through 1.2.39, from 1.3.0 through 1.3.6, from 2.0.0 through 2.0.13. Users are recommended to upgrade to version Tomcat Native 1.3.7 or 2.0.14 and Tomcat 11.0.20, 10.1.53 and 9.0.116, which fix the issue.
CVE-2026-31958	Tornado is a Python web framework and asynchronous networking library. In versions of Tornado prior to 6.5.5, the only limit on the number of parts in multipart/form-data is the max_body_size setting (default 100MB). Since parsing occurs synchronously on the main thread, this creates the possibility of denial-of-service due to the cost of parsing very large multipart bodies with many parts. This vulnerability is fixed in 6.5.5.
CVE-2026-32597	PyJWT is a JSON Web Token implementation in Python. Prior to 2.12.0, PyJWT does not validate the crit (Critical) Header Parameter defined in RFC 7515 §4.1.11. When a JWS token contains a crit array listing extensions that PyJWT does not understand, the library accepts the token instead of rejecting it. This violates the MUST requirement in the RFC. This vulnerability is fixed in 2.12.0.
CVE-2026-32990	Improper Input Validation vulnerability in Apache Tomcat due to an incomplete fix of CVE-2025-66614. This issue affects Apache Tomcat: from 11.0.15 through 11.0.19, from 10.1.50 through 10.1.52, from 9.0.113 through 9.0.115. Users are recommended to upgrade to version 11.0.20, 10.1.53 or 9.0.116, which fix the issue.
CVE-2026-34481	Apache Log4j's JsonTemplateLayout https://logging.apache.org/log4j/2.x/manual/json-template-layout.html , in versions up to and including 2.25.3, produces invalid JSON output when log events contain non-finite floating-point values (NaN, Infinity, or -Infinity), which are prohibited by RFC 8259. This may cause downstream log processing systems to reject or fail to index affected records. An attacker can exploit this issue only if both of the following conditions are met: * The application uses JsonTemplateLayout. * The application logs a MapMessage containing an attacker-controlled floating-point value. Users are advised to upgrade to Apache Log4j JSON Template Layout 2.25.4, which corrects this issue.

CVE ID	Description
CVE-2026-34483	Improper Encoding or Escaping of Output vulnerability in the JsonAccessLogValve component of Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.20, from 10.1.0-M1 through 10.1.53, from 9.0.40 through 9.0.116. Users are recommended to upgrade to version 11.0.21, 10.1.54 or 9.0.117 , which fix the issue.
CVE-2026-34487	Insertion of Sensitive Information into Log File vulnerability in the cloud membership for clustering component of Apache Tomcat exposed the Kubernetes bearer token. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.20, from 10.1.0-M1 through 10.1.53, from 9.0.13 through 9.0.116. Users are recommended to upgrade to version 11.0.21, 10.1.54 or 9.0.117, which fix the issue.
CVE-2026-34500	CLIENT_CERT authentication does not fail as expected for some scenarios when soft fail is disabled and FFM is used in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M14 through 11.0.20, from 10.1.22 through 10.1.53, from 9.0.92 through 9.0.116. Users are recommended to upgrade to version 11.0.21, 10.1.54 or 9.0.117, which fixes the issue.
CVE-2026-35536	In Tornado before 6.5.5, cookie attribute injection could occur because the domain, path, and samesite arguments to .RequestHandler.set_cookie were not checked for crafted characters.
CVE-2026-41284	Allocation of Resources Without Limits or Throttling vulnerability in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.21, from 10.1.0-M1 through 10.1.54, from 9.0.0.M1 through 9.0.117. Older, unsupported versions may also be affected. Users are recommended to upgrade to version [FIXED_VERSION], which fixes the issue.
CVE-2026-41293	Improper Input Validation vulnerability in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.21, from 10.1.0-M1 through 10.1.54, from 9.0.0.M1 through 9.0.117, from 10.0.0-M1 through 10.0.27. Older, end of support versions may also be affected. Users are recommended to upgrade to version [FIXED_VERSION], which fixes the issue.
CVE-2026-42304	Twisted is an event-based framework for internet applications, supporting Python 3.6+. Prior to 26.4.0rc2, the twisted.names module is vulnerable to a Denial of Service (DoS) attack via resource exhaustion during DNS name decompression. A remote, unauthenticated attacker can exploit this by sending a crafted TCP DNS packet containing deeply chained compression pointers. This flaw bypasses previous loop-prevention logic, causing the single-threaded Twisted reactor to hang while processing millions of recursive lookups, effectively freezing the server. This vulnerability is fixed in 26.4.0rc2.
CVE-2026-42498	Exposure of HTTP Authentication Header to unexpected hosts during WebSocket authentication vulnerability in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.21, from 10.1.0-M1 through 10.1.54, from 9.0.2 through 9.0.117, from 8.5.24 through 8.5.100, from 7.0.83 through 7.0.109. Users are recommended to upgrade to version 11.0.22, 10.1.55 or 9.0.118, which fix the issue.
CVE-2026-42577	Netty is an asynchronous, event-driven network application framework. From 4.2.0.Final to 4.2.13.Final , Netty's epoll transport fails to detect and close TCP connections that receive a RST after being half-closed, leading to stale channels that are never cleaned up and, in some code paths, a 100% CPU busy-loop in the event loop thread. This vulnerability is fixed in 4.2.13.Final.
CVE-2026-43512	DEPRECATED: Authentication Bypass Issues vulnerability in digest authentication in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.21, from 10.1.0-M1 through 10.1.54, from 9.0.0.M1 through 9.0.117, from 8.5.0 through 8.5.100, from before 7.0.0. Older unsupported versions any also be affect Users are recommended to upgrade to version 11.0.22, 10.1.55 or 9.0.118 which fix the issue.
CVE-2026-43513	Improper Handling of Case Sensitivity vulnerability in LockOutRealm in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.21, from 10.1.0-M1 through 10.1.54, from 9.0.0.M1 through 9.0.117, from 8.5.0 through 8.5.100, from 7.0.0 through 7.0.109. Older unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.22, 10.1.55 or 9.0.118 which fix the issue.
CVE-2026-43514	Observable Timing Discrepancy vulnerability when comparing AJP secret in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.21, from 10.1.0-M1 through 10.1.54, from 9.0.0.M1 through 9.0.117, from 8.5.0 through 8.5.100, from 7.0.0 through 7.0.109. Older unsupported versions may also be affected. Users are recommended to upgrade to version 11.0.22, 10.1.55 or 9.0.118 which fix the issue.
CVE-2026-43515	Improper Authorization vulnerability when multiple method constraints define an HTTP method for the same extension in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.21, from 10.1.0-M1 through 10.1.54, from 9.0.0.M1 through 9.0.117, from 8.5.0 through 8.5.100, from 7.0.0 through 7.0.109. Users are recommended to upgrade to version 11.0.22, 10.1.55 or 9.0.118 which fix the issue.
CVE-2026-44660	UltraJSON is a fast JSON encoder and decoder written in pure C with bindings for Python 3.7+. Prior to 5.12.1, when ujson.dump() writes to a file-like object and the write operation raises an exception, the serialized JSON string object is not decremented, leaking memory. Each failed write operation leaks the full size of the serialized payload. This vulnerability is fixed in 5.12.1.

CVE ID	Description
CVE-2026-48522	PyJWT is a JSON Web Token implementation in Python. Prior to 2.13.0, PyJWKClient passes its uri argument directly to urllib.request.urlopen() which uses Python stdlib's default OpenerDirector registering HTTPHandler, HTTPSHandler, FTPHandler, FileHandler, and DataHandler. There is currently no documented option to restrict which schemes PyJWKClient will fetch. If an application's jku URL ingestion path accepts attacker-influenced URLs (e.g., from JWT header, configuration file, OAuth flow parameter), the attacker can cause PyJWKClient to read arbitrary local files via file:// (SSRF on local filesystem), cause PyJWKClient to attempt FTP / data-URI fetches (broader SSRF surface), or forge tokens that PyJWT verifies as valid. The library does not directly return non-HTTP(S) URI contents to the attacker; the chained "plant a JWKS to forge tokens" scenario described in the original report requires additional application-layer flaws (attacker write access to a filesystem path, untrusted jku derivation) that this fix does not address. This vulnerability is fixed in 2.13.0.
CVE-2026-48524	PyJWT is a JSON Web Token implementation in Python. Prior to 2.13.0, PyJWKClient.get_signing_key() forces a fresh HTTP request to the JWKS endpoint for every JWT with an unknown kid value, with no rate limiting. Since kid comes from the unverified token header, an attacker can trigger unlimited outbound requests. The vulnerability surfaces only when a JWKS fetch fails; an attacker can attempt to provoke that with sustained unknown-kid traffic, but the outcome depends on upstream JWKS-endpoint behavior (rate limiting, transient errors) which is beyond the attacker's control. This vulnerability is fixed in 2.13.0.
CVE-2026-48526	PyJWT is a JSON Web Token implementation in Python. Prior to 2.13.0, when the verifier is decoding JSON Web Tokens, while supporting both asymmetric and HMAC algorithms, the library does not validate use of JSON Web Keys in HMAC algorithm, allowing attacker to use the issuer public key as the secret key for HMAC algorithm. This vulnerability is fixed in 2.13.0.
CVE-2026-49853	Tornado is a Python web framework and asynchronous networking library. Prior to 6.5.6, SimpleAsyncHTTPClient shallow-copied redirected requests and removed only the Host header, leaving Authorization, auth_username, auth_password, and auth_mode in place when a redirect changed scheme, host, or port. This issue is fixed in version 6.5.6.
CVE-2026-49854	Tornado is a Python web framework and asynchronous networking library. Prior to 6.5.6, the optional native extension tornado.speedups implemented websocket_mask without validating that the mask argument is exactly four bytes, allowing the C function to read up to three bytes beyond the provided buffer when reached through Tornado XSRF token decoding with the native extension active. This issue is fixed in version 6.5.6.
CVE-2026-49855	Tornado is a Python web framework and asynchronous networking library. Prior to 6.5.6, Tornado gzip decompression routines processed limited-size chunks but did not enforce an overall limit on accumulated decompressed chunks, allowing a malicious server accessed by SimpleAsyncHTTPClient or an HTTPServer configured with decompress_request=True to consume effectively unlimited memory. This issue is fixed in version 6.5.6.
CVE-2026-50193	jackson-databind contains the general-purpose data-binding functionality and tree-model for Jackson Data Processor. From 2.13.0 until 2.14.0, a potential Denial-of-Service exists when attacker sends deeply nested JSON if (and only if) the service reads deeply nested (1000s of levels) JSON as JsonNode (ObjectMapper.readTree()) and writes out same (or modified) node using JsonNode.toString(). This can consume significant amount of resources with concurrent relatively small requests (1000 nested arrays is 2kB). This vulnerability is fixed in 2.14.0.
CVE-2026-54911	UltraJSON is a fast JSON encoder and decoder written in pure C with bindings for Python 3.7+. Prior to 5.13.0, ujson.dumps() (or ujson.dump() or ujson.encode()) have a reject_bytes=False option. When set, they may accept malformed or truncated UTF-8 byte sequences, silently rewriting them into different Unicode characters instead of rejecting them. This leads to input validation bypass and data integrity issues. This vulnerability is fixed in 5.13.0
CVE-2026-55760	Handlebars.java provides logic-less and semantic Mustache templates with Java. Prior to 4.5.2, applications that pass user-controlled input to Handlebars.compile() using FileTemplateLoader or ClassPathTemplateLoader are vulnerable to path traversal, allowing arbitrary file read through template names derived from URL path parameters, request parameters, or other user-controlled sources. This issue is fixed in version 4.5.2.

Cloudera Runtime 7.3.2

Common vulnerabilities and Exposures (CVEs) fixed in Cloudera Runtime 7.3.2.

Table 16: CVE IDs and description

CVE ID	Description
CVE-2016-1000027	Pivotal Spring Framework through 5.3.16 suffers from a potential remote code execution (RCE) issue if used for Java deserialization of untrusted data. Depending on how the library is implemented within a product, this issue may or not occur, and authentication may be required. NOTE: the vendor's position is that untrusted data is not an intended use case. The product's behavior will not be changed because some users rely on deserialization of trusted data.

CVE ID	Description
CVE-2021-22570	Nullptr dereference when a null char is present in a proto symbol. The symbol is parsed incorrectly, leading to an unchecked call into the proto file's name during generation of the resulting error message. Since the symbol is incorrectly parsed, the file is nullptr. We recommend upgrading to version 3.15.0 or greater.
CVE-2021-22569	An issue in protobuf-java allowed the interleaving of com.google.protobuf.UnknownFieldSet fields in such a way that would be processed out of order. A small malicious payload can occupy the parser for several minutes by creating large numbers of short-lived objects that cause frequent, repeated pauses. We recommend upgrading libraries beyond the vulnerable versions.
CVE-2024-29133	Out-of-bounds Write vulnerability in Apache Commons Configuration.This issue affects Apache Commons Configuration: from 2.0 before 2.10.1. Users are recommended to upgrade to version 2.10.1, which fixes the issue.
CVE-2024-6763	Eclipse Jetty is a lightweight, highly scalable, Java-based web server and Servlet engine . It includes a utility class, HttpURI, for URI/URL parsing. The HttpURI class does insufficient validation on the authority segment of a URI. However the behaviour of HttpURI differs from the common browsers in how it handles a URI that would be considered invalid if fully validated against the RRC. Specifically HttpURI and the browser may differ on the value of the host extracted from an invalid URI and thus a combination of Jetty and a vulnerable browser may be vulnerable to a open redirect attack or to a SSRF attack if the URI is used after passing validation checks.
CVE-2023-33201	Bouncy Castle For Java before 1.74 is affected by an LDAP injection vulnerability. The vulnerability only affects applications that use an LDAP CertStore from Bouncy Castle to validate X.509 certificates. During the certificate validation process, Bouncy Castle inserts the certificate's Subject Name into an LDAP search filter without any escaping, which leads to an LDAP injection vulnerability.
CVE-2024-30171	An issue was discovered in Bouncy Castle Java TLS API and JSSE Provider before 1.78. Timing-based leakage may occur in RSA based handshakes because of exception processing
CVE-2023-42503	Improper Input Validation, Uncontrolled Resource Consumption vulnerability in Apache Commons Compress in TAR parsing.This issue affects Apache Commons Compress: from 1.22 before 1.24.0. Users are recommended to upgrade to version 1.24.0, which fixes the issue. A third party can create a malformed TAR file by manipulating file modification times headers, which when parsed with Apache Commons Compress, will cause a denial of service issue via CPU consumption. In version 1.22 of Apache Commons Compress, support was added for file modification times with higher precision (issue # COMPRESS-612 [1]). The format for the PAX extended headers carrying this data consists of two numbers separated by a period [2], indicating seconds and subsecond precision (for example “1647221103.5998539”). The impacted fields are “atime”, “ctime”, “mtime” and “LIBARCHIVE.creationtime”. No input validation is performed prior to the parsing of header values. Parsing of these numbers uses the BigDecimal [3] class from the JDK which has a publicly known algorithmic complexity issue when doing operations on large numbers, causing denial of service (see issue # JDK-6560193 [4]). A third party can manipulate file time headers in a TAR file by placing a number with a very long fraction (300,000 digits) or a number with exponent notation (such as “9e9999999”) within a file modification time header, and the parsing of files with these headers will take hours instead of seconds, leading to a denial of service via exhaustion of CPU resources. This issue is similar to CVE-2012-2098 [5]. [1]: https://issues.apache.org/jira/browse/COMPRESS-612 [2]: https://pubs.opengroup.org/onlinepubs/9699919799/utilities/pax.html#tag_20_92_13_05 [3]: https://docs.oracle.com/javase/8/docs/api/java/math/BigDecimal.html [4]: https://bugs.openjdk.org/browse/JDK-6560193 [5]: https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2012-2098 Only applications using CompressorStreamFactory class (with auto-detection of file types), TarArchiveInputStream and TarFile classes to parse TAR files are impacted. Since this code was introduced in v1.22, only that version and later versions are impacted.
CVE-2021-4048	An out-of-bounds read flaw was found in the CLARRV, DLARRV, SLARRV, and ZLARRV functions in lapack through version 3.10.0, as also used in OpenBLAS before version 0.3.18. Specially crafted inputs passed to these functions could cause an application using lapack to crash or possibly disclose portions of its memory.
CVE-2024-38821	Spring WebFlux applications that have Spring Security authorization rules on static resources can be bypassed under certain circumstances. For this to impact an application, all of the following must be true: * It must be a WebFlux application * It must be using Spring's static resources support * It must have a non-permitAll authorization rule applied to the static resources support

CVE ID	Description
CVE-2024-52046	The ObjectSerializationDecoder in Apache MINA uses Java's native deserialization protocol to process incoming serialized data but lacks the necessary security checks and defenses. This vulnerability allows attackers to exploit the deserialization process by sending specially crafted malicious serialized data, potentially leading to remote code execution (RCE) attacks. This issue affects MINA core versions 2.0.X, 2.1.X and 2.2.X, and will be fixed by the releases 2.0.27, 2.1.10 and 2.2.4. It's also important to note that an application using MINA core library will only be affected if the IoBuffer#getObject() method is called, and this specific method is potentially called when adding a ProtocolCodecFilter instance using the ObjectSerializationCodecFactory class in the filter chain. If your application is specifically using those classes, you have to upgrade to the latest version of MINA core library. Upgrading will not be enough: you also need to explicitly allow the classes the decoder will accept in the ObjectSerializationDecoder instance, using one of the three new methods: /** * Accept class names where the supplied ClassNameMatcher matches for * deserialization, unless they are otherwise rejected. * * @param classNameMatcher the matcher to use */ public void accept(ClassNameMatcher classNameMatcher) /** * Accept class names that match the supplied pattern for * deserialization, unless they are otherwise rejected. * * @param pattern standard Java regexp */ public void accept(Pattern pattern) /** * Accept the wildcard specified classes for deserialization, * unless they are otherwise rejected. * * @param patterns Wildcard file name patterns as defined by * {@link org.apache.commons.io.FilenameUtils#wildcardMatch(String, String) FilenameUtils.wildcardMatch} */ public void accept(String... patterns) By default, the decoder will reject *all* classes that will be present in the incoming data. Note: The FtpServer, SSHd and Vysper sub-project are not affected by this issue.
CVE-2024-52316	Unchecked Error Condition vulnerability in Apache Tomcat. If Tomcat is configured to use a custom Jakarta Authentication (formerly JASPIC) ServerAuthContext component which may throw an exception during the authentication process without explicitly setting an HTTP status to indicate failure, the authentication may not fail, allowing the user to bypass the authentication process. There are no known Jakarta Authentication components that behave in this way. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.0-M26, from 10.1.0-M1 through 10.1.30, from 9.0.0-M1 through 9.0.95. Users are recommended to upgrade to version 11.0.0, 10.1.31 or 9.0.96, which fix the issue.
CVE-2024-53990	The AsyncHttpClient (AHC) library allows Java applications to easily execute HTTP requests and asynchronously process HTTP responses. When making any HTTP request, the automatically enabled and self-managed CookieStore (aka cookie jar) will silently replace explicitly defined Cookies with any that have the same name from the cookie jar. For services that operate with multiple users, this can result in one user's Cookie being used for another user's requests.
CVE-2025-24813	Path Equivalence: 'file.Name' (Internal Dot) leading to Remote Code Execution and/or Information disclosure and/or malicious content added to uploaded files via write enabled Default Servlet in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.2, from 10.1.0-M1 through 10.1.34, from 9.0.0.M1 through 9.0.98. The following versions were EOL at the time the CVE was created but are known to be affected: 8.5.0 though 8.5.100. Other, older, EOL versions may also be affected. If all of the following were true, a malicious user was able to view security sensitive files and/or inject content into those files: - writes enabled for the default servlet (disabled by default) - support for partial PUT (enabled by default) - a target URL for security sensitive uploads that was a sub-directory of a target URL for public uploads - attacker knowledge of the names of security sensitive files being uploaded - the security sensitive files also being uploaded via partial PUT If all of the following were true, a malicious user was able to perform remote code execution: - writes enabled for the default servlet (disabled by default) - support for partial PUT (enabled by default) - application was using Tomcat's file based session persistence with the default storage location - application included a library that may be leveraged in a deserialization attack Users are recommended to upgrade to version 11.0.3, 10.1.35 or 9.0.99, which fixes the issue.
CVE-2025-31651	Improper Neutralization of Escape, Meta, or Control Sequences vulnerability in Apache Tomcat. For a subset of unlikely rewrite rule configurations, it was possible for a specially crafted request to bypass some rewrite rules. If those rewrite rules effectively enforced security constraints, those constraints could be bypassed. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.5, from 10.1.0-M1 through 10.1.39, from 9.0.0.M1 through 9.0.102. The following versions were EOL at the time the CVE was created but are known to be affected: 8.5.0 though 8.5.100. Other, older, EOL versions may also be affected. Users are recommended to upgrade to version [FIXED_VERSION], which fixes the issue.
CVE-2024-48910	DOMPurify is a DOM-only, super-fast, uber-tolerant XSS sanitizer for HTML, MathML and SVG. DOMPurify was vulnerable to prototype pollution. This vulnerability is fixed in 2.4.2.
CVE-2024-45216	Improper Authentication vulnerability in Apache Solr. Solr instances using the PKIAuthenticationPlugin, which is enabled by default when Solr Authentication is used, are vulnerable to Authentication bypass. A fake ending at the end of any Solr API URL path, will allow requests to skip Authentication while maintaining the API contract with the original URL Path. This fake ending looks like an unprotected API path, however it is stripped off internally after authentication but before API routing. This issue affects Apache Solr: from 5.3.0 before 8.11.4, from 9.0.0 before 9.7.0. Users are recommended to upgrade to version 9.7.0, or 8.11.4, which fix the issue.

CVE ID	Description
CVE-2026-27727	mchange-commons-java, a library that provides Java utilities, includes code that mirrors early implementations of JNDI functionality, including support for remote `factoryClassLocation` values, by which code can be downloaded and invoked within a running application. If an attacker can provoke an application to read a maliciously crafted `jaxax.naming.Reference` or serialized object, they can provoke the download and execution of malicious code. Implementations of this functionality within the JDK were disabled by default behind a System property that defaults to `false`, `com.sun.jndi.ldap.object.trustURLCodebase`. However, since mchange-commons-java includes an independent implementation of JNDI derefencing, libraries (such as c3p0) that resolve references via that implementation could be provoked to download and execute malicious code even after the JDK was hardened. Mirroring the JDK patch, mchange-commons-java's JNDI functionality is gated by configuration parameters that default to restrictive values starting in version 0.4.0. No known workarounds are available. Versions prior to 0.4.0 should be avoided on application CLASSPATHs.
CVE-2025-66614	Improper Input Validation vulnerability. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.14, from 10.1.0-M1 through 10.1.49, from 9.0.0-M1 through 9.0.112. The following versions were EOL at the time the CVE was created but are known to be affected: 8.5.0 through 8.5.100. Older EOL versions are not affected. Tomcat did not validate that the host name provided via the SNI extension was the same as the host name provided in the HTTP host header field. If Tomcat was configured with more than one virtual host and the TLS configuration for one of those hosts did not require client certificate authentication but another one did, it was possible for a client to bypass the client certificate authentication by sending different host names in the SNI extension and the HTTP host header field. The vulnerability only applies if client certificate authentication is only enforced at the Connector. It does not apply if client certificate authentication is enforced at the web application. Users are recommended to upgrade to version 11.0.15 or later, 10.1.50 or later or 9.0.113 or later, which fix the issue.
CVE-2025-24813	Path Equivalence: 'file.Name' (Internal Dot) leading to Remote Code Execution and/or Information disclosure and/or malicious content added to uploaded files via write enabled Default Servlet in Apache Tomcat. This issue affects Apache Tomcat: from 11.0.0-M1 through 11.0.2, from 10.1.0-M1 through 10.1.34, from 9.0.0.M1 through 9.0.98. The following versions were EOL at the time the CVE was created but are known to be affected: 8.5.0 though 8.5.100. Other, older, EOL versions may also be affected. If all of the following were true, a malicious user was able to view security sensitive files and/or inject content into those files: - writes enabled for the default servlet (disabled by default) - support for partial PUT (enabled by default) - a target URL for security sensitive uploads that was a sub-directory of a target URL for public uploads - attacker knowledge of the names of security sensitive files being uploaded - the security sensitive files also being uploaded via partial PUT If all of the following were true, a malicious user was able to perform remote code execution: - writes enabled for the default servlet (disabled by default) - support for partial PUT (enabled by default) - application was using Tomcat's file based session persistence with the default storage location - application included a library that may be leveraged in a deserialization attack Users are recommended to upgrade to version 11.0.3, 10.1.35 or 9.0.99, which fixes the issue.
CVE-2017-1000028	Oracle, GlassFish Server Open Source Edition 4.1 is vulnerable to both authenticated and unauthenticated Directory Traversal vulnerability, that can be exploited by issuing a specially crafted HTTP GET request.
CVE-2019-18218	cdf_read_property_info in cdf.c in file through 5.37 does not restrict the number of CDF_VECTOR elements, which allows a heap-based buffer overflow (4-byte out-of-bounds write).
CVE-2021-0341	In verifyHostName of OkHostnameVerifier.java, there is a possible way to accept a certificate for the wrong domain due to improperly used crypto. This could lead to remote information disclosure with no additional execution privileges needed. User interaction is not needed for exploitation.Product: AndroidVersions: Android-8.1 Android-9 Android-10 Android-11Android ID: A-171980069
CVE-2021-44878	If an OpenID Connect provider supports the "none" algorithm (i.e., tokens with no signature), pac4j v5.3.0 (and prior) does not refuse it without an explicit configuration on its side or for the "idtoken" response type which is not secure and violates the OpenID Core Specification. The "none" algorithm does not require any signature verification when validating the ID tokens, which allows the attacker to bypass the token validation by injecting a malformed ID token using "none" as the value of "alg" key in the header with an empty signature value.
CVE-2022-25844	The package angular after 1.7.0 are vulnerable to Regular Expression Denial of Service (ReDoS) by providing a custom locale rule that makes it possible to assign the parameter in posPre: ' '.repeat() of NUMBER_FORMATS.PATTERNS[1].posPre with a very high value. **Note:** 1) This package has been deprecated and is no longer maintained. 2) The vulnerable versions are 1.7.0 and higher.
CVE-2022-3171	A parsing issue with binary data in protobuf-java core and lite versions prior to 3.21.7, 3.20.3, 3.19.6 and 3.16.3 can lead to a denial of service attack. Inputs containing multiple instances of non-repeated embedded messages with repeated or unknown fields causes objects to be converted back-n-forth between mutable and immutable forms, resulting in potentially long garbage collection pauses. We recommend updating to the versions mentioned above.
CVE-2022-3509	A parsing issue similar to CVE-2022-3171, but with textformat in protobuf-java core and lite versions prior to 3.21.7, 3.20.3, 3.19.6 and 3.16.3 can lead to a denial of service attack. Inputs containing multiple instances of non-repeated embedded messages with repeated or unknown fields causes objects to be converted back-n-forth between mutable and immutable forms, resulting in potentially long garbage collection pauses. We recommend updating to the versions mentioned above.

CVE ID	Description
CVE-2022-3510	A parsing issue similar to CVE-2022-3171, but with Message-Type Extensions in protobuf-java core and lite versions prior to 3.21.7, 3.20.3, 3.19.6 and 3.16.3 can lead to a denial of service attack. Inputs containing multiple instances of non-repeated embedded messages with repeated or unknown fields causes objects to be converted back-n-forth between mutable and immutable forms, resulting in potentially long garbage collection pauses. We recommend updating to the versions mentioned above.
CVE-2023-3635	GzipSource does not handle an exception that might be raised when parsing a malformed gzip buffer. This may lead to denial of service of the Okio client when handling a crafted GZIP archive, by using the GzipSource class.
CVE-2023-39410	When deserializing untrusted or corrupted data, it is possible for a reader to consume memory beyond the allowed constraints and thus lead to out of memory on the system. This issue affects Java applications using Apache Avro Java SDK up to and including 1.11.2. Users should update to apache-avro version 1.11.3 which addresses this issue.
CVE-2023-4586	A vulnerability was found in the Hot Rod client. This security issue occurs as the Hot Rod client does not enable hostname validation when using TLS, possibly resulting in a man-in-the-middle (MITM) attack.
CVE-2023-7272	In Eclipse Parsson before 1.0.4 and 1.1.3, a document with a large depth of nested objects can allow an attacker to cause a Java stack overflow exception and denial of service. Eclipse Parsson allows processing (e.g. parse, generate, transform and query) JSON documents.
CVE-2024-13009	In Eclipse Jetty versions 9.4.0 to 9.4.56 a buffer can be incorrectly released when confronted with a gzip error when inflating a request body. This can result in corrupted and/or inadvertent sharing of data between requests.
CVE-2024-21490	This affects versions of the package angular from 1.3.0. A regular expression used to split the value of the ng-srcset directive is vulnerable to super-linear runtime due to backtracking. With large carefully-crafted input, this can result in catastrophic backtracking and cause a denial of service. **Note:** This package is EOL and will not receive any updates to address this issue. Users should migrate to [@angular/core](https://www.npmjs.com/package/@angular/core).
CVE-2024-25638	dnsjava is an implementation of DNS in Java. Records in DNS replies are not checked for their relevance to the query, allowing an attacker to respond with RRs from different zones. This vulnerability is fixed in 3.6.0.
CVE-2024-29131	Out-of-bounds Write vulnerability in Apache Commons Configuration.This issue affects Apache Commons Configuration: from 2.0 before 2.10.1. Users are recommended to upgrade to version 2.10.1, which fixes the issue.
CVE-2024-29857	An issue was discovered in ECCurve.java and ECCurve.cs in Bouncy Castle Java (BC Java) before 1.78, BC Java LTS before 2.73.6, BC-FJA before 1.0.2.5, and BC C# .Net before 2.3.1. Importing an EC certificate with crafted F2m parameters can lead to excessive CPU consumption during the evaluation of the curve parameters.
CVE-2024-34447	An issue was discovered in the Bouncy Castle Crypto Package For Java before BC TLS Java 1.0.19 (ships with BC Java 1.78, BC Java (LTS) 2.73.6) and before BC FIPS TLS Java 1.0.19. When endpoint identification is enabled in the BCJSSE and an SSL socket is created without an explicit hostname (as happens with HttpsURLConnection), hostname verification could be performed against a DNS-resolved IP address in some situations, opening up a possibility of DNS poisoning.
CVE-2024-47561	Schema parsing in the Java SDK of Apache Avro 1.11.3 and previous versions allows bad actors to execute arbitrary code. Users are recommended to upgrade to version 1.11.4 or 1.12.0, which fix this issue.
CVE-2024-57699	A security issue was found in Netplex Json-smart 2.5.0 through 2.5.1. When loading a specially crafted JSON input, containing a large number of '{', a stack exhaustion can be trigger, which could allow an attacker to cause a Denial of Service (DoS). This issue exists because of an incomplete fix for CVE-2023-1370.
CVE-2024-7254	Any project that parses untrusted Protocol Buffers data containing an arbitrary number of nested groups / series of SGROUP tags can corrupted by exceeding the stack limit i.e. StackOverflow. Parsing nested groups as unknown fields with DiscardUnknownFieldsParser or Java Protobuf Lite parser, or against Protobuf map fields, creates unbounded recursions that can be abused by an attacker.
CVE-2025-1948	In Eclipse Jetty versions 12.0.0 to 12.0.16 included, an HTTP/2 client can specify a very large value for the HTTP/2 settings parameter SETTINGS_MAX_HEADER_LIST_SIZE. The Jetty HTTP/2 server does not perform validation on this setting, and tries to allocate a ByteBuffer of the specified capacity to encode HTTP responses, likely resulting in OutOfMemoryError being thrown, or even the JVM process exiting.
CVE-2025-22228	BCryptPasswordEncoder.matches(CharSequence,String) will incorrectly return true for passwords larger than 72 characters as long as the first 72 characters are the same.
CVE-2025-24970	Netty, an asynchronous, event-driven network application framework, has a vulnerability starting in version 4.1.91.Final and prior to version 4.1.118.Final. When a special crafted packet is received via SslHandler it doesn't correctly handle validation of such a packet in all cases which can lead to a native crash. Version 4.1.118.Final contains a patch. As workaround its possible to either disable the usage of the native SSLEngine or change the code manually.

CVE ID	Description
CVE-2025-31650	Improper Input Validation vulnerability in Apache Tomcat. Incorrect error handling for some invalid HTTP priority headers resulted in incomplete clean-up of the failed request which created a memory leak. A large number of such requests could trigger an OutOfMemoryException resulting in a denial of service. This issue affects Apache Tomcat: from 9.0.76 through 9.0.102, from 10.1.10 through 10.1.39, from 11.0.0-M2 through 11.0.5. The following versions were EOL at the time the CVE was created but are known to be affected: 8.5.90 though 8.5.100. Users are recommended to upgrade to version 9.0.104, 10.1.40 or 11.0.6 which fix the issue.
CVE-2025-46762	Schema parsing in the parquet-avro module of Apache Parquet 1.15.0 and previous versions allows bad actors to execute arbitrary code. While 1.15.1 introduced a fix to restrict untrusted packages, the default setting of trusted packages still allows malicious classes from these packages to be executed. The exploit is only applicable if the client code of parquet-avro uses the "specific" or the "reflect" models deliberately for reading Parquet files. ("generic" model is not impacted) Users are recommended to upgrade to 1.15.2 or set the system property "org.apache.parquet.avro.SERIALIZABLE_PACKAGES" to an empty string on 1.15.1. Both are sufficient to fix the issue.
CVE-2025-67721	Aircompressor is a library with ports of the Snappy, LZO, LZ4, and Zstandard compression algorithms to Java. In versions 3.3 and below, incorrect handling of malformed data in Java-based decompressor implementations for Snappy and LZ4 allow remote attackers to read previous buffer contents via crafted compressed input. With certain crafted compressed inputs, elements from the output buffer can end up in the uncompressed output, potentially leaking sensitive data. This is relevant for applications that reuse the same output buffer to uncompress multiple inputs. This can be the case of a web server that allocates a fix-sized buffer for performance purposes. There is similar vulnerability in GHSA-cmp6-m4wj-q63q. This issue is fixed in version 3.4.
CVE-2025-55163	Netty is an asynchronous, event-driven network application framework. Prior to versions 4.1.124.Final and 4.2.4.Final, Netty is vulnerable to MadeYouReset DDoS. This is a logical vulnerability in the HTTP/2 protocol, that uses malformed HTTP/2 control frames in order to break the max concurrent streams limit - which results in resource exhaustion and distributed denial of service. This issue has been patched in versions 4.1.124.Final and 4.2.4.Final.
CVE-2025-58057	Netty is an asynchronous event-driven network application framework for rapid development of maintainable high performance protocol servers & clients. In netty-codec-compression versions 4.1.124.Final and below, and netty-codec versions 4.2.4.Final and below, when supplied with specially crafted input, BrotliDecoder and certain other decompression decoders will allocate a large number of reachable byte buffers, which can lead to denial of service. BrotliDecoder.decompress has no limit in how often it calls pull, decompressing data 64K bytes at a time. The buffers are saved in the output list, and remain reachable until OOM is hit. This is fixed in versions 4.1.125.Final of netty-codec and 4.2.5.Final of netty-codec-compression.
CVE-2025-58056	Netty is an asynchronous event-driven network application framework for development of maintainable high performance protocol servers and clients. In versions 4.1.124.Final, and 4.2.0.Alpha3 through 4.2.4.Final, Netty incorrectly accepts standalone newline characters (LF) as a chunk-size line terminator, regardless of a preceding carriage return (CR), instead of requiring CRLF per HTTP/1.1 standards. When combined with reverse proxies that parse LF differently (treating it as part of the chunk extension), attackers can craft requests that the proxy sees as one request but Netty processes as two, enabling request smuggling attacks. This is fixed in versions 4.1.125.Final and 4.2.5.Final.
CVE-2025-48734	Improper Access Control vulnerability in Apache Commons. A special BeanIntrospector class was added in version 1.9.2. This can be used to stop attackers from using the declared class property of Java enum objects to get access to the classloader. However this protection was not enabled by default. PropertyUtilsBean (and consequently BeanUtilsBean) now disallows declared class level property access by default. Releases 1.11.0 and 2.0.0-M2 address a potential security issue when accessing enum properties in an uncontrolled way. If an application using Commons BeanUtils passes property paths from an external source directly to the getProperty() method of PropertyUtilsBean, an attacker can access the enum's class loader via the "declaredClass" property available on all Java "enum" objects. Accessing the enum's "declaredClass" allows remote attackers to access the ClassLoader and execute arbitrary code. The same issue exists with PropertyUtilsBean.getNestedProperty(). Starting in versions 1.11.0 and 2.0.0-M2 a special BeanIntrospector suppresses the "declaredClass" property. Note that this new BeanIntrospector is enabled by default, but you can disable it to regain the old behavior; see section 2.5 of the user's guide and the unit tests. This issue affects Apache Commons BeanUtils 1.x before 1.11.0, and 2.x before 2.0.0-M2.Users of the artifact commons-beanutils:commons-beanutils 1.x are recommended to upgrade to version 1.11.0, which fixes the issue. Users of the artifact org.apache.commons:commons-beanutils2 2.x are recommended to upgrade to version 2.0.0-M2, which fixes the issue.
CVE-2026-24308	Improper handling of configuration values in ZKConfig in Apache ZooKeeper 3.8.5 and 3.9.4 on all platforms allows an attacker to expose sensitive information stored in client configuration in the client's logfile. Configuration values are exposed at INFO level logging rendering potential production systems affected by the issue. Users are recommended to upgrade to version 3.8.6 or 3.9.5 which fixes this issue.
CVE-2026-24281	Hostname verification in Apache ZooKeeper ZKTrustManager falls back to reverse DNS (PTR) when IP SAN validation fails, allowing attackers who control or spoof PTR records to impersonate ZooKeeper servers or clients with a valid certificate for the PTR name. It's important to note that attacker must present a certificate which is trusted by ZKTrustManager which makes the attack vector harder to exploit. Users are recommended to upgrade to version 3.8.6 or 3.9.5, which fixes this issue by introducing a new configuration option to disable reverse DNS lookup in client and quorum protocols.

CVE ID	Description
CVE-2023-1428	There exists an vulnerability causing an abort() to be called in gRPC. The following headers cause gRPC's C++ implementation to abort() when called via http2: te: x (x != trailers) :scheme: x (x != http, https) grpclb_client_stats: x (x == anything) On top of sending one of those headers, a later header must be sent that gets the total header size past 8KB. We recommend upgrading past git commit 2485fa94bd8a723e5c977d55a3ce10b301b437f8 or v1.53 and above.
CVE-2026-22022	Deployments of Apache Solr 5.3.0 through 9.10.0 that rely on Solr's "Rule Based Authorization Plugin" are vulnerable to allowing unauthorized access to certain Solr APIs, due to insufficiently strict input validation in those components. Only deployments that meet all of the following criteria are impacted by this vulnerability: * Use of Solr's "RuleBasedAuthorizationPlugin" * A RuleBasedAuthorizationPlugin config (see security.json) that specifies multiple "roles" * A RuleBasedAuthorizationPlugin permission list (see security.json) that uses one or more of the following pre-defined permission rules: "config-read", "config-edit", "schema-read", "metrics-read", or "security-read". * A RuleBasedAuthorizationPlugin permission list that doesn't define the "all" pre-defined permission * A networking setup that allows clients to make unfiltered network requests to Solr. (i.e. user-submitted HTTP/HTTPS requests reach Solr as-is, unmodified or restricted by any intervening proxy or gateway) Users can mitigate this vulnerability by ensuring that their RuleBasedAuthorizationPlugin configuration specifies the "all" pre-defined permission and associates the permission with an "admin" or other privileged role. Users can also upgrade to a Solr version outside of the impacted range, such as the recently released Solr 9.10.1.
CVE-2024-45217	Insecure Default Initialization of Resource vulnerability in Apache Solr. New ConfigSets that are created via a Restore command, which copy a configSet from the backup and give it a new name, are created without setting the "trusted" metadata. ConfigSets that do not contain the flag are trusted implicitly if the metadata is missing, therefore this leads to "trusted" ConfigSets that may not have been created with an Authenticated request. "trusted" ConfigSets are able to load custom code into classloaders, therefore the flag is supposed to only be set when the request that uploads the ConfigSet is Authenticated & Authorized. This issue affects Apache Solr: from 6.6.0 before 8.11.4, from 9.0.0 before 9.7.0. This issue does not affect Solr instances that are secured via Authentication/Authorization. Users are primarily recommended to use Authentication and Authorization when running Solr. However, upgrading to version 9.7.0, or 8.11.4 will mitigate this issue otherwise.
CVE-2026-22444	The "create core" API of Apache Solr 8.6 through 9.10.0 lacks sufficient input validation on some API parameters, which can cause Solr to check the existence of and attempt to read file-system paths that should be disallowed by Solr's "allowPaths" security setting https://https://solr.apache.org/guide/solr/latest/configuration-guide/configuring-solr-xml.html#the-solr-element . These read-only accesses can allow users to create cores using unexpected configsets if any are accessible via the filesystem. On Windows systems configured to allow UNC paths this can additionally cause disclosure of NTLM "user" hashes. Solr deployments are subject to this vulnerability if they meet the following criteria: * Solr is running in its "standalone" mode. * Solr's "allowPath" setting is being used to restrict file access to certain directories. * Solr's "create core" API is exposed and accessible to untrusted users. This can happen if Solr's RuleBasedAuthorizationPlugin https://solr.apache.org/guide/solr/latest/deployment-guide/rule-based-authorization-plugin.html is disabled, or if it is enabled but the "core-admin-edit" predefined permission (or an equivalent custom permission) is given to low-trust (i.e. non-admin) user roles. Users can mitigate this by enabling Solr's RuleBasedAuthorizationPlugin (if disabled) and configuring a permission-list that prevents untrusted users from creating new Solr cores. Users should also upgrade to Apache Solr 9.10.1 or greater, which contain fixes for this issue.
CVE-2025-5222	A stack buffer overflow was found in Internationl components for unicode (ICU). While running the genrb binary, the 'subtag' struct overflowed at the SRBRoot::addTag function. This issue may lead to memory corruption and local arbitrary code execution.
CVE-2025-5115	In Eclipse Jetty, versions <=9.4.57, <=10.0.25, <=11.0.25, <=12.0.21, <=12.1.0.alpha2, an HTTP/2 client may trigger the server to send RST_STREAM frames, for example by sending frames that are malformed or that should not be sent in a particular stream state, therefore forcing the server to consume resources such as CPU and memory. For example, a client can open a stream and then send WINDOW_UPDATE frames with window size increment of 0, which is illegal. Per specification https://www.rfc-editor.org/rfc/rfc9113.html#name-window_update , the server should send a RST_STREAM frame. The client can now open another stream and send another bad WINDOW_UPDATE, therefore causing the server to consume more resources than necessary, as this case does not exceed the max number of concurrent streams, yet the client is able to create an enormous amount of streams in a short period of time. The attack can be performed with other conditions (for example, a DATA frame for a closed stream) that cause the server to send a RST_STREAM frame. Links: * https://github.com/jetty/jetty.project/security/advisories/GHSA-mmxm-8w33-wc4h
CVE-2019-10768	In AngularJS before 1.7.9 the function `merge()` could be tricked into adding or modifying properties of `Object.prototype` using a `__proto__` payload.
CVE-2022-2048	In Eclipse Jetty HTTP/2 server implementation, when encountering an invalid HTTP/2 request, the error handling has a bug that can wind up not properly cleaning up the active connections and associated resources. This can lead to a Denial of Service scenario where there are no enough resources left to process good requests.

CVE ID	Description
CVE-2023-36478	Eclipse Jetty provides a web server and servlet container. In versions 11.0.0 through 11.0.15, 10.0.0 through 10.0.15, and 9.0.0 through 9.4.52, an integer overflow in `MetaDataBuilder.checkSize` allows for HTTP/2 HPACK header values to exceed their size limit. `MetaDataBuilder.java` determines if a header name or value exceeds the size limit, and throws an exception if the limit is exceeded. However, when length is very large and huffman is true, the multiplication by 4 in line 295 will overflow, and length will become negative. `(_size+length)` will now be negative, and the check on line 296 will not be triggered. Furthermore, `MetaDataBuilder.checkSize` allows for user-entered HPACK header value sizes to be negative, potentially leading to a very large buffer allocation later on when the user-entered size is multiplied by 2. This means that if a user provides a negative length value (or, more precisely, a length value which, when multiplied by the 4/3 fudge factor, is negative), and this length value is a very large positive number when multiplied by 2, then the user can cause a very large buffer to be allocated on the server. Users of HTTP/2 can be impacted by a remote denial of service attack. The issue has been fixed in versions 11.0.16, 10.0.16, and 9.4.53. There are no known workarounds.
CVE-2023-44487	The HTTP/2 protocol allows a denial of service (server resource consumption) because request cancellation can reset many streams quickly, as exploited in the wild in August through October 2023.
CVE-2024-22201	Jetty is a Java based web server and servlet engine. An HTTP/2 SSL connection that is established and TCP congested will be leaked when it times out. An attacker can cause many connections to end up in this state, and the server may run out of file descriptors, eventually causing the server to stop accepting new connections from valid clients. The vulnerability is patched in 9.4.54, 10.0.20, 11.0.20, and 12.0.6.
CVE-2024-9823	There exists a security vulnerability in Jetty's DosFilter which can be exploited by unauthorized users to cause remote denial-of-service (DoS) attack on the server using DosFilter. By repeatedly sending crafted requests, attackers can trigger OutofMemory errors and exhaust the server's memory finally.
CVE-2023-50386	Improper Control of Dynamically-Managed Code Resources, Unrestricted Upload of File with Dangerous Type, Inclusion of Functionality from Untrusted Control Sphere vulnerability in Apache Solr.This issue affects Apache Solr: from 6.0.0 through 8.11.2, from 9.0.0 before 9.4.1. In the affected versions, Solr ConfigSets accepted Java jar and class files to be uploaded through the ConfigSets API. When backing up Solr Collections, these configSet files would be saved to disk when using the LocalFileSystemRepository (the default for backups). If the backup was saved to a directory that Solr uses in its ClassPath/ClassLoaders, then the jar and class files would be available to use with any ConfigSet, trusted or untrusted. When Solr is run in a secure way (Authorization enabled), as is strongly suggested, this vulnerability is limited to extending the Backup permissions with the ability to add libraries. Users are recommended to upgrade to version 8.11.3 or 9.4.1, which fix the issue. In these versions, the following protections have been added: * Users are no longer able to upload files to a configSet that could be executed via a Java ClassLoader. * The Backup API restricts saving backups to directories that are used in the ClassLoader.
CVE-2023-50291	Insufficiently Protected Credentials vulnerability in Apache Solr. This issue affects Apache Solr: from 6.0.0 through 8.11.2, from 9.0.0 before 9.3.0. One of the two endpoints that publishes the Solr process' Java system properties, /admin/info/properties, was only setup to hide system properties that had "password" contained in the name. There are a number of sensitive system properties, such as "basicauth" and "aws.secretKey" do not contain "password", thus their values were published via the "/admin/info/properties" endpoint. This endpoint populates the list of System Properties on the home screen of the Solr Admin page, making the exposed credentials visible in the UI. This /admin/info/properties endpoint is protected under the "config-read" permission. Therefore, Solr Clouds with Authorization enabled will only be vulnerable through logged-in users that have the "config-read" permission. Users are recommended to upgrade to version 9.3.0 or 8.11.3, which fixes the issue. A single option now controls hiding Java system property for all endpoints, "-Dsolr.hiddenSysProps". By default all known sensitive properties are hidden (including "-Dbasicauth"), as well as any property with a name containing "secret" or "password". Users who cannot upgrade can also use the following Java system property to fix the issue: '-Dsolr.redaction.system.pattern=.*(password|secret|basicauth).*'
CVE-2023-50292	Incorrect Permission Assignment for Critical Resource, Improper Control of Dynamically-Managed Code Resources vulnerability in Apache Solr. This issue affects Apache Solr: from 8.10.0 through 8.11.2, from 9.0.0 before 9.3.0. The Schema Designer was introduced to allow users to more easily configure and test new Schemas and configSets. However, when the feature was created, the "trust" (authentication) of these configSets was not considered. External library loading is only available to configSets that are "trusted" (created by authenticated users), thus non-authenticated users are unable to perform Remote Code Execution. Since the Schema Designer loaded configSets without taking their "trust" into account, configSets that were created by unauthenticated users were allowed to load external libraries when used in the Schema Designer. Users are recommended to upgrade to version 9.3.0, which fixes the issue.
CVE-2023-50298	Exposure of Sensitive Information to an Unauthorized Actor vulnerability in Apache Solr.This issue affects Apache Solr: from 6.0.0 through 8.11.2, from 9.0.0 before 9.4.1. Solr Streaming Expressions allows users to extract data from other Solr Clouds, using a "zkHost" parameter. When original SolrCloud is setup to use ZooKeeper credentials and ACLs, they will be sent to whatever "zkHost" the user provides. An attacker could setup a server to mock ZooKeeper, that accepts ZooKeeper requests with credentials and ACLs and extracts the sensitive information, then send a streaming expression using the mock server's address in "zkHost". Streaming Expressions are exposed via the "/streaming" handler, with "read" permissions. Users are recommended to upgrade to version 8.11.3 or 9.4.1, which fix the issue. From these versions on, only zkHost values that have the same server address (regardless of chroot), will use the given ZooKeeper credentials and ACLs when connecting.

CVE ID	Description
CVE-2025-66524	Apache NiFi 1.20.0 through 2.6.0 include the GetAsanaObject Processor, which requires integration with a configurable Distribute Map Cache Client Service for storing and retrieving state information. The GetAsanaObject Processor used generic Java Object serialization and deserialization without filtering. Unfiltered Java object deserialization does not provide protection against crafted state information stored in the cache server configured for GetAsanaObject. Exploitation requires an Apache NiFi system running with the GetAsanaObject Processor, and direct access to the configured cache server. Upgrading to Apache NiFi 2.7.0 is the recommended mitigation, which replaces Java Object serialization with JSON serialization. Removing the GetAsanaObject Processor located in the nifi-asana-processors-nar bundle also prevents exploitation.
CVE-2020-14734	Vulnerability in the Oracle Text component of Oracle Database Server. Supported versions that are affected are 11.2.0.4, 12.1.0.2, 12.2.0.1, 18c and 19c. Difficult to exploit vulnerability allows unauthenticated attacker with network access via Oracle Net to compromise Oracle Text. Successful attacks of this vulnerability can result in takeover of Oracle Text. CVSS 3.1 Base Score 8.1 (Confidentiality, Integrity and Availability impacts). CVSS Vector: (CVSS:3.1/AV:N/AC:H/PR:N/UI:N/S:U/C:H/I:H/A:H).
CVE-2017-15288	The compilation daemon in Scala before 2.10.7, 2.11.x before 2.11.12, and 2.12.x before 2.12.4 uses weak permissions for private files in /tmp/scala-devel/${USER:shared}/scalac-compile-server-port, which allows local users to write to arbitrary class files and consequently gain privileges.
CVE-2023-45142	OpenTelemetry-Go Contrib is a collection of third-party packages for OpenTelemetry-Go. A handler wrapper out of the box adds labels `http.user_agent` and `http.method` that have unbound cardinality. It leads to the server's potential memory exhaustion when many malicious requests are sent to it. HTTP header User-Agent or HTTP method for requests can be easily set by an attacker to be random and long. The library internally uses `httpconv.ServerRequest` that records every value for HTTP `method` and `User-Agent`. In order to be affected, a program has to use the `otelhttp.NewHandler` wrapper and not filter any unknown HTTP methods or User agents on the level of CDN, LB, previous middleware, etc. Version 0.44.0 fixed this issue when the values collected for attribute `http.request.method` were changed to be restricted to a set of well-known values and other high cardinality attributes were removed. As a workaround to stop being affected, `otelhttp.WithFilter()` can be used, but it requires manual careful configuration to not log certain requests entirely. For convenience and safe usage of this library, it should by default mark with the label `unknown` non-standard HTTP methods and User agents to show that such requests were made but do not increase cardinality. In case someone wants to stay with the current behavior, library API should allow to enable it.
CVE-2023-47108	OpenTelemetry-Go Contrib is a collection of third-party packages for OpenTelemetry-Go. Starting in version 0.37.0 and prior to version 0.46.0, the grpc Unary Server Interceptor out of the box adds labels `net.peer.sock.addr` and `net.peer.sock.port` that have unbound cardinality. It leads to the server's potential memory exhaustion when many malicious requests are sent. An attacker can easily flood the peer address and port for requests. Version 0.46.0 contains a fix for this issue. As a workaround to stop being affected, a view removing the attributes can be used. The other possibility is to disable grpc metrics instrumentation by passing `otelgrpc.WithMeterProvider` option with `noop.NewMeterProvider`.
CVE-2025-54920	This issue affects Apache Spark: before 3.5.7 and 4.0.1. Users are recommended to upgrade to version 3.5.7 or 4.0.1 and above, which fixes the issue. Summary Apache Spark 3.5.4 and earlier versions contain a code execution vulnerability in the Spark History Web UI due to overly permissive Jackson deserialization of event log data. This allows an attacker with access to the Spark event logs directory to inject malicious JSON payloads that trigger deserialization of arbitrary classes, enabling command execution on the host running the Spark History Server. Details The vulnerability arises because the Spark History Server uses Jackson polymorphic deserialization with @JsonTypeInfo.Id.CLASS on SparkListenerEvent objects, allowing an attacker to specify arbitrary class names in the event JSON. This behavior permits instantiating unintended classes, such as org.apache.hive.jdbc.HiveConnection, which can perform network calls or other malicious actions during deserialization. The attacker can exploit this by injecting crafted JSON content into the Spark event log files, which the History Server then deserializes on startup or when loading event logs. For example, the attacker can force the History Server to open a JDBC connection to a remote attacker-controlled server, demonstrating remote command injection capability. Proof of Concept: 1. Run Spark with event logging enabled, writing to a writable directory (spark-logs). 2. Inject the following JSON at the beginning of an event log file: { "Event": "org.apache.hive.jdbc.HiveConnection", "uri": "jdbc:hive2://<IP>:<PORT>/", "info": { "hive.metastore.uris": "thrift://<IP>:<PORT>" } } 3. Start the Spark History Server with logs pointing to the modified directory. 4. The Spark History Server initiates a JDBC connection to the attacker’s server, confirming the injection. Impact An attacker with write access to Spark event logs can execute arbitrary code on the server running the History Server, potentially compromising the entire system.
CVE-2025-54988	Critical XXE in Apache Tika (tika-parser-pdf-module) in Apache Tika 1.13 through and including 3.2.1 on all platforms allows an attacker to carry out XML External Entity injection via a crafted XFA file inside of a PDF. An attacker may be able to read sensitive data or trigger malicious requests to internal resources or third-party servers. Note that the tika-parser-pdf-module is used as a dependency in several Tika packages including at least: tika-parsers-standard-modules, tika-parsers-standard-package, tika-app, tika-grpc and tika-server-standard. Users are recommended to upgrade to version 3.2.2, which fixes this issue.

CVE ID	Description
CVE-2026-24734	Improper Input Validation vulnerability in Apache Tomcat Native, Apache Tomcat. When using an OCSP responder, Tomcat Native (and Tomcat's FFM port of the Tomcat Native code) did not complete verification or freshness checks on the OCSP response which could allow certificate revocation to be bypassed. This issue affects Apache Tomcat Native: from 1.3.0 through 1.3.4, from 2.0.0 through 2.0.11; Apache Tomcat: from 11.0.0-M1 through 11.0.17, from 10.1.0-M7 through 10.1.51, from 9.0.83 through 9.0.114. The following versions were EOL at the time the CVE was created but are known to be affected: from 1.1.23 through 1.1.34, from 1.2.0 through 1.2.39. Older EOL versions are not affected. Apache Tomcat Native users are recommended to upgrade to versions 1.3.5 or later or 2.0.12 or later, which fix the issue. Apache Tomcat users are recommended to upgrade to versions 11.0.18 or later, 10.1.52 or later or 9.0.115 or later which fix the issue.
CVE-2026-25639	Axios is a promise based HTTP client for the browser and Node.js. Prior to versions 0.30.3 and 1.13.5, the mergeConfig function in axios crashes with a TypeError when processing configuration objects containing __proto__ as an own property. An attacker can trigger this by providing a malicious configuration object created via JSON.parse(), causing complete denial of service. This vulnerability is fixed in versions 0.30.3 and 1.13.5.
CVE-2025-11965	In Eclipse Vert.x versions [4.0.0, 4.5.21] and [5.0.0, 5.0.4], a StaticHandler configuration for restricting access to hidden files fails to restrict access to hidden directories, allowing unauthorized users to retrieve files within them (e.g. '.git/config').
CVE-2017-18214	The moment module before 2.19.3 for Node.js is prone to a regular expression denial of service via a crafted date string, a different vulnerability than CVE-2016-4055.
CVE-2022-24785	Moment.js is a JavaScript date library for parsing, validating, manipulating, and formatting dates. A path traversal vulnerability impacts npm (server) users of Moment.js between versions 1.0.1 and 2.29.1, especially if a user-provided locale string is directly used to switch moment locale. This problem is patched in 2.29.2, and the patch can be applied to all affected versions. As a workaround, sanitize the user-provided locale name before passing it to Moment.js.
CVE-2022-1271	An arbitrary file write vulnerability was found in GNU gzip's zgrep utility. When zgrep is applied on the attacker's chosen file name (for example, a crafted file name), this can overwrite an attacker's content to an arbitrary attacker-selected file. This flaw occurs due to insufficient validation when processing filenames with two or more newlines where selected content and the target file names are embedded in crafted multi-line file names. This flaw allows a remote, low privileged attacker to force zgrep to write arbitrary files on the system.
CVE-2015-4035	scripts/xzgrep.in in xzgrep 5.2.x before 5.2.0, before 5.0.0 does not properly process file names containing semicolons, which allows remote attackers to execute arbitrary code by having a user run xzgrep on a crafted file name.

Cloudera products has a transitive dependencies on one of the AWS library that is reporting these vulnerabilities. However, Cloudera products do not use this library in any of the vulnerable code paths as noted in the CVE. As such Cloudera product remains not vulnerable to this. When fixes from AWS becomes available, our products will upgrade to the latest versions in upcoming releases. The following CVEs are not vulnerable:

CVE ID	Description
CVE-2025-55163	Netty is an asynchronous, event-driven network application framework. Prior to versions 4.1.124.Final and 4.2.4.Final, Netty is vulnerable to MadeYouReset DDoS. This is a logical vulnerability in the HTTP/2 protocol, that uses malformed HTTP/2 control frames in order to break the max concurrent streams limit - which results in resource exhaustion and distributed denial of service. This issue has been patched in versions 4.1.124.Final and 4.2.4.Final.
CVE-2025-24970	Netty, an asynchronous, event-driven network application framework, has a vulnerability starting in version 4.1.91.Final and prior to version 4.1.118.Final. When a special crafted packet is received via SslHandler it doesn't correctly handle validation of such a packet in all cases which can lead to a native crash. Version 4.1.118.Final contains a patch. As workaround its possible to either disable the usage of the native SSLEngine or change the code manually.
CVE-2025-58057	Netty is an asynchronous event-driven network application framework for rapid development of maintainable high performance protocol servers & clients. In netty-codec-compression versions 4.1.124.Final and below, and netty-codec versions 4.2.4.Final and below, when supplied with specially crafted input, BrotliDecoder and certain other decompression decoders will allocate a large number of reachable byte buffers, which can lead to denial of service. BrotliDecoder.decompress has no limit in how often it calls pull, decompressing data 64K bytes at a time. The buffers are saved in the output list, and remain reachable until OOM is hit. This is fixed in versions 4.1.125.Final of netty-codec and 4.2.5.Final of netty-codec-compression.
CVE-2024-47535	Netty is an asynchronous event-driven network application framework for rapid development of maintainable high performance protocol servers & clients. An unsafe reading of environment file could potentially cause a denial of service in Netty. When loaded on an Windows application, Netty attempts to load a file that does not exist. If an attacker creates such a large file, the Netty application crashes. This vulnerability is fixed in 4.1.115.

CVE ID	Description
CVE-2025-25193	Netty, an asynchronous, event-driven network application framework, has a vulnerability in versions up to and including 4.1.118.Final. An unsafe reading of environment file could potentially cause a denial of service in Netty. When loaded on an Windows application, Netty attempts to load a file that does not exist. If an attacker creates such a large file, the Netty application crash. A similar issue was previously reported as CVE-2024-47535. This issue was fixed, but the fix was incomplete in that null-bytes were not counted against the input limit. Commit d1fbda62d3a47835d3fb35db8bd42ecc205a5386 contains an updated fix.
CVE-2025-58056	Netty is an asynchronous event-driven network application framework for development of maintainable high performance protocol servers and clients. In versions 4.1.124.Final, and 4.2.0.Alpha3 through 4.2.4.Final, Netty incorrectly accepts standalone newline characters (LF) as a chunk-size line terminator, regardless of a preceding carriage return (CR), instead of requiring CRLF per HTTP/1.1 standards. When combined with reverse proxies that parse LF differently (treating it as part of the chunk extension), attackers can craft requests that the proxy sees as one request but Netty processes as two, enabling request smuggling attacks. This is fixed in versions 4.1.125.Final and 4.2.5.Final.

Deprecation Notices In Cloudera Runtime 7.3.2

Components and features that will be deprecated or removed in this release or a future release.

Terminology

Items in this section are designated as follows:

Deprecated

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

Moving

Technology that Cloudera is moving from a future Cloudera 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 Cloudera release and plan for the alternative Cloudera offering or subscription for the technology.

Removed

Technology that Cloudera has removed from Cloudera 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

- Apache Spark 2

  Spark 3 is the default Spark version in Cloudera Runtime. Spark 2 (and Livy 2) has been removed and no longer available in 7.3.1

  

  Important:

  Spark 3 contains a large number of changes from Spark 2.

  Refer to *Upgrading Spark* for more information on upgrading Spark clusters to 7.3.1.0, and *Migrating Spark Applications* for more information on migrating your existing Spark applications between versions 2 and 3.
- Apache Livy 2 (see Deprecation Notices for Apache Livy)
- Apache Zeppelin (see Deprecation Notices for Apache Zeppelin)

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

Platform and OS

Cloudera Runtime 7.3.2 removes support for several operating systems, databases, Python versions, and Java Development Kits (JDKs). To ensure a successful upgrade, verify that your environment meets these new requirements.

Cloudera Runtime 7.3.2 SP1 (7.3.2.10000)

Database Support

Starting with Cloudera Runtime 7.3.2 SP1, Cloudera deprecates the following databases, which are no longer recommended for use:

- PostgreSQL 14
- Including MariaDB 10.6 and all lower 10.x release series (such as 10.5 and 10.4)
- Including MariaDB 11.3 and all lower 11.x release series (such as 11.2 and 11.1)
- MySQL 8

Operating System

Starting with Cloudera Runtime 7.3.2 SP1, Cloudera deprecates the following operating systems, which are no longer recommended for use:

- Including RHEL or Oracle Linux version 9.5 and lower
- RHEL 8.10 with FIPS mode enabled
- RHEL 8.8 with FIPS mode enabled
- RHEL or Oracle Linux 8.8
- RHEL 8.9



Important: FIPS 140-2 support continues until September 21, 2026.

Cloudera Runtime 7.3.2

Database Support

Starting with Cloudera Runtime 7.3.2, Clouderaremoves support for the following databases, which are no longer supported:

- PostgreSQL 13
- Oracle 19c (including 19c RAC)

Operating system

Starting with Cloudera Runtime 7.3.2, Cloudera removes support for the following operating systems, which are no longer supported:

- SLES 15 SP4
- Ubuntu 20.04

Python Support

Cloudera Runtime 7.3.2 limits Python compatibility to modern versions, and now only supports Python 3.11. Consequently, Cloudera has removed support for:

- Python 3.8
- Python 3.9
- Python 3.10

JDK support

To improve performance and security, Cloudera Runtime 7.3.2 consolidates Java support, and now only supports JDK 17. Consequently, Cloudera has removed support for:

- Azul JDK 8, Open JDK 8, and Oracle JDK 8

- sAzul JDK 11, Open JDK 11, and Oracle JDK 11

Deprecation Notices for Cruise Control

Certain features and functionality in Cruise Control are deprecated or removed in Cloudera Runtime 7.3.2.

Removed

Cloudera Manager metrics reporter

Cloudera Manager metrics reporter is removed from Cloudera Runtime 7.3.2. Migrate to Cruise Control metrics reporter.

Deprecation Notices for Apache Impala

Python 2 support for the impala-shell command is deprecated in Cloudera Runtime 7.3.2.10000 SP1. Run impala-shell with Python 3, which is the only supported version in current and future releases.

Deprecation Notices for Apache Kafka

Certain features and functionality in Apache Kafka are deprecated or removed in Cloudera Runtime 7.3.2. You must review these changes along with the information about the features in Kafka that will be removed or deprecated in a future release.

Important: The following list of deprecated and removed items is not exhaustive and only contains items that have a direct and immediate effect on Kafka in Cloudera. For a full list of deprecation and/or removals in the version Apache Kafka shipped with Cloudera Runtime, review the *Notable Changes* as well as the *Release Notes* on https://kafka.apache.org/.

Deprecated

ZooKeeper

Deploying new or using existing Kafka clusters running in ZooKeeper mode is deprecated. Support for ZooKeeper-based Kafka clusters will be removed in a future release. Cloudera recommends the following:

- Deploy all new Kafka clusters in KRaft mode.
- Migrate existing ZooKeeper-based clusters to KRaft following an upgrade to Cloudera Runtime 7.3.2.

  This is the only version where migration is possible. Neither previous or future major, minor, and maintenance versions support migration.

  For additional information, see the following resources:

  - Installing Kafka in Cloudera Base on premises
  - Kafka KRaft overview
  - Migrating Kafka from ZooKeeper to KRaft overview

MirrorMaker (MM1)

MirrorMaker is deprecated. Cloudera recommends that you use Streams Replication Manager instead.

--zookeeper

The --zookeeper option is only supported for the kafka-configs tool and should be only used when updating SCRAM Credential configurations. The --zookeeper option is either deprecated in or

removed from other Kafka command line tools. Cloudera recommends that you use the --bootstrap-server option instead.

Deprecation Notices for Apache Livy

Certain features and functionality in Apache Livy are deprecated or removed in Cloudera Runtime 7.3.1. You must review these changes along with the information about the features in Livy that will be removed or deprecated in a future release.

Removed

Apache Livy 2

As Spark 3 is the default Spark version in Cloudera Runtime, Livy 2 has been removed, alongside with Spark 2, and no longer available in 7.3.1



Important:

Spark 3 contains a large number of changes from Spark 2.

For more information on upgrading to Spark 3, refer to Upgrading Spark for more information on upgrading Spark clusters to 7.3.1, and Migrating Spark Applications for more information on migrating your existing Spark applications between versions 2 and 3.

Deprecation Notices for Apache Ozone

Certain features and functionality in Apache Ozone are deprecated or removed in Cloudera Runtime 7.3.2.0.

Removed

Ordered deletion of snapshots

Starting from Cloudera Runtime 7.1.9 SP1 CHF12 onwards, ordered deletion is removed after significant fixes and improvements to Snapshot Garbage Collection.

Configurations related to ordered deletion ozone.snapshot.ordered.deletion.enabled and ozone.snapshot.deep.cleaning.enabled are irrelevant and have been removed from Cloudera Manager.

Deprecation Notices for Apache Ranger

Certain features and functionality in Apache Ranger are deprecated or removed in Cloudera Runtime 7.3.2.0.

Removed

Deprecation of Ranger Backbone Classic UI

Starting from Cloudera Runtime 7.3.2.0, the Backbone Classic UI option has been removed from the Ranger Admin interface. Support for the legacy Backbone-based UI is no longer available.

Deprecation of CSV and Excel policy exports

Exporting policies in CSV and Excel formats from the Ranger Admin UI Reports page is deprecated due to security risks associated with CSV injection.

Deprecation Notices for Apache Spark

Certain features and functionality in Apache Spark 2 are deprecated or removed in Cloudera Runtime 7.3.1. You must review these changes along with the information about the features in Spark 2 that will be removed or deprecated in a future release.

Removed

Apache Spark 2

Spark 3 is the default Spark version in Cloudera Runtime. Spark 2 has been removed and no longer available in 7.3.1



Important:

Spark 3 contains a large number of changes from Spark 2.

Refer to Upgrading Spark for more information on upgrading Spark clusters to 7.3.1, and Migrating Spark Applications for more information on migrating your existing Spark applications between versions 2 and 3.

Deprecation Notices for Streams Replication Manager

Certain features and functionality in Streams Replication Manager are deprecated or removed in Cloudera Runtime 7.3.2.

Deprecated

Legacy per-flow REST server implementation

The legacy per-flow REST server implementation of Streams Replication Manager is deprecated in Cloudera Runtime 7.3.2 and will be removed in a future release. In the previous implementation, the Streams Replication Manager Driver role started a dedicated REST server for each replication flow. The new implementation uses a single REST server for all flows. Cloudera recommends that you migrate your Streams Replication Manager clusters to the new implementation as part of upgrading to 7.3.2. For migration steps, see Configuring the SRM REST server for replication.

Deprecation Notices for Apache Zeppelin

Certain features and functionality in Apache Zeppelin are deprecated or removed in Cloudera Runtime 7.3.1. You must review these changes along with the information about the features in Zeppelin that will be removed or deprecated in a future release.

Removed

Apache Zeppelin

Apache Zeppelin is removed from Cloudera on premises.

You can reinstall the Zeppelin service as an external CSD and restore and use old notebooks, but Cloudera does not provide for Zeppelin starting in 7.3.1 and above. For more information, see Reinstall Apache Zeppelin in 7.3.1

Cloudera recommends you back up all existing Zeppelin notebooks before upgrading to version 7.3.1.

API Modifications and Removals

All the APIs that have either been modified or removed for this release.

Components with API differences in Cloudera Runtime 7.3.2.10000 SP1

The following is a list of all the APIs that have been modified or removed in this release of Cloudera Runtime. The APIs listed here are on a component by component basis and may include recommendations for how to manage the change.

API Compatibility changes in 7.3.2 for HBase

Removed or Modified APIs in Cloudera Runtime 7.3.2 for HBase and recommendations for how to handle them.

Apache Version of HBase in 7.3.2 and 7.3.2.10000 is 2.6.3.

Modified APIs in 7.3.2

The following HBase API endpoints have been modified in Cloudera Runtime 7.3.2.

RegionMetrics

Interface Modified

Package Name

org.apache.hadoop.hbase

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method getCurrentRegionColdDataRatio() in org.apache.hadoop.hbase.RegionMetrics.

Reason for change

HBASE-30102: Add metric to account for region data classified as cold by the Time Based Priority logic

Recommendation

Recompile client application.

Recompilation Required?

Yes

API Compatibility changes in 7.3.2 for ZooKeeper

Removed or Modified APIs in Cloudera Runtime 7.3.2 for ZooKeeper and recommendations for how to handle them.

Apache Version of ZooKeeper in 7.3.2 and 7.3.2.10000 is 3.8.5

Removed APIs in 7.3.2

The following ZooKeeper API endpoints have been removed from Cloudera Runtime 7.3.2.

SecurityUtils.createSaslServer

Method Removed

Package Name

org.apache.zookeeper.util

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

ZOOKEEPER-5043: Disable fallback to DIGEST-MD5 in SaslServer when FIPS mode is set

Recommendation

Recompile client application. The method now has a new parameter: ZKConfig config.

Recompilation Required?

Yes

X509Util.createSSLContextAndOptionsFromConfig

Method Removed

Package Name

org.apache.zookeeper.common

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

CDPD-104916: ZooKeeper client fails to connect on FIPS-enabled TLS 1.3 clusters due to TLS protocol mismatch

Recommendation

Recompile client application. The method now has a new parameter: boolean isClient.

Recompilation Required?

Yes

X509Util.createTrustManager

Method Removed

Package Name

org.apache.zookeeper.common

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

ZOOKEEPER-4986: Disable reverse DNS lookup in TLS client and server

Recommendation

Recompile client application. The method now has a new parameter: final boolean allowReverseDnsLookup.

Recompilation Required?

Yes

Modified APIs in 7.3.2

The following ZooKeeper API endpoints have been modified in Cloudera Runtime 7.3.2.

X509Util

Interface/Class Modified

Package Name

org.apache.zookeeper.common

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method shouldAllowReverseDnsLookup() in org.apache.zookeeper.common.X509Util.

Reason for change

ZOOKEEPER-4986: Disable reverse DNS lookup in TLS client and server

Recommendation

Recompile client application. Abstract method boolean shouldAllowReverseDnsLookup() has been added to this class.

Recompilation Required?

Yes

X509Util.DEFAULT_PROTOCOL

Field Renamed

Package Name

org.apache.zookeeper.common

Effect

Recompilation of a client program may be terminated with the message: cannot find variable DEFAULT_PROTOCOL in org.apache.zookeeper.common.X509Util.

Reason for change

CDPD-104916: ZooKeeper client fails to connect on FIPS-enabled TLS 1.3 clusters due to TLS protocol mismatch

Recommendation

Recompile client application. Field DEFAULT_PROTOCOL has been renamed to DEFAULT_TLS_PROTOCOL_CLIENT.

Recompilation Required?

Yes

PrometheusMetricsProvider.getServlet

Method Modified

Package Name

org.apache.zookeeper.metrics.prometheus

Effect

Recompilation of a client program may be terminated with the message: incompatible types: io.prometheus.metrics.exporter.servlet.javax.PrometheusMetricsServlet cannot be converted to org.apache.zookeeper.metrics.prometheus.PrometheusMetricsProvider.MetricsServletImpl.

Reason for change

ZOOKEEPER-4741: High latency under heavy load when prometheus metrics enabled.

Recommendation

Recompile client application. Return value type has been changed from org.apache.zookeeper.metrics.prometheus.PrometheusMetricsProvider.MetricsServletImpl to io.prometheus.metrics.exporter.servlet.javax.PrometheusMetricsServlet.

Recompilation Required?

Yes

Components with API differences in Cloudera Runtime 7.3.2.100 CHF 1

The following is a list of all the APIs that have been modified or removed in this release of Cloudera Runtime. The APIs listed here are on a component by component basis and may include recommendations for how to manage the change.

API Compatibility changes in 7.3.2.100 for ZooKeeper

Removed or Modified APIs in Cloudera Runtime 7.3.2.100 CHF 1 for ZooKeeper and recommendations for how to handle them.

Apache Version of ZooKeeper in 7.3.2 was 3.8.5 and Apache Version of ZooKeeper in 7.3.2.100 is 3.8.5

Removed APIs in 7.3.2

The following APIs are no longer available for ZooKeeper in Cloudera Runtime 7.3.2.100.

X509Util.createTrustManager

Method Added

Package Name

org.apache.zookeeper.common

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

ZOOKEEPER-4986

Recommendation

New method parameter was added: boolean allowReverseDnsLookup

Recompilation Required?

Yes

Components with API differences in Cloudera Runtime 7.3.2

The following is a list of all the APIs that have been modified or removed in this release of Cloudera Runtime. The APIs listed here are on a component by component basis and may include recommendations for how to manage the change.

API Compatibility changes in 7.3.2 for Hadoop

Removed or Modified APIs in Cloudera Runtime 7.3.2 for Hadoop and recommendations for how to handle them.

Apache Version of Hadoop in 7.3.1 was 3.1.1 and Apache Version of Hadoop in 7.3.2 is 3.4.2

Removed APIs in 7.3.2

The following Hadoop API endpoints have been removed from Cloudera Runtime 7.3.2. This change is a result of the fix for CVE-2025-48924 and YARN's migration to common-lang 3.

GetApplicationsRequest.getFinishRange

Method Removed

Package Name

org.apache.hadoop.yarn.api.protocolrecords

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

YARN-10772

Recommendation

Now returns Range<Long> instead of LongRange.

Recompilation Required?

Yes

GetApplicationsRequest.getStartRange

Method Removed

Package Name

org.apache.hadoop.yarn.api.protocolrecords

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

YARN-10772

Recommendation

Now returns Range<Long> instead of LongRange.

Recompilation Required?

Yes

GetApplicationsRequest.newInstance

Method Removed

Package Name

org.apache.hadoop.yarn.api.protocolrecords

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

YARN-10772

Recommendation

factory method parameter types changed

Recompilation Required?

Yes

GetApplicationsRequest.setFinishRange

Method Removed

Package Name

org.apache.hadoop.yarn.api.protocolrecords

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

YARN-10772

Recommendation

parameter type changed: request.setFinishRange(new LongRange(startTime, endTime)) uses commons-lang. Now chnaged to request.setFinishRange(Range.between(startTime, endTime)); // commons-lang3

Recompilation Required?

Yes

GetApplicationsRequest.setStartRange

Method Removed

Package Name

org.apache.hadoop.yarn.api.protocolrecords

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

YARN-10772

Recommendation

parameter type changed: request.setStartRange(new LongRange(startTime, endTime)) uses commons-lang. Now changed to request.setStartRange(Range.between(startTime, endTime)); // commons-lang3

Recompilation Required?

Yes

API Compatibility changes in 7.3.2 for HBase

Removed or Modified APIs in Cloudera Runtime 7.3.2 for HBase and recommendations for how to handle them.

Apache Version of HBase in 7.3.1.600 was 2.4.17 and Apache Version of HBase in 7.3.2 is 2.6.3.

Modified APIs in 7.3.2

The following APIs have been modified for HBase in Cloudera Runtime 7.3.2

DataBlockEncoding.createEncoder

Access level has been changed from protected to package-private.

Package Name

org.apache.hadoop.hbase.io.encoding

Effect

Recompilation of a client program may be terminated with the message: `createEncoder (String)` has package-private access in `org.apache.hadoop.hbase.io.encoding.DataBlockEncoding`.

Reason for change

HBASE-27206

Recommendation

Use `DataBlockEncoding.getEncoder` instead of `createEncoder`.

Recompilation Required?

YES

Bytes.toByteString

Return value type has been changed from `com.google.protobuf.ByteString` to `org.apache.hbase.thirdparty.com.google.protobuf.ByteString`.

Package Name

org.apache.hadoop.hbase.util

Effect

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

Reason for change

CDPD-33849

Recommendation

Method signature changed, return type changed to HBase thirdparty protobuf class. If a client application calls this method, it should import the HBase thirdparty protobuf class.

Recompilation Required?

YES

PBType.inputStreamFromByteRange

Return value type has been changed from `com.google.protobuf.CodedInputStream` to `org.apache.hbase.thirdparty.com.google.protobuf.CodedInputStream`.

Package Name

org.apache.hadoop.hbase.types

Effect

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

Reason for change

CDPD-33849

Recommendation

Method signature changed, return type changed to HBase thirdparty protobuf class. If a client application calls this method, it should import the HBase thirdparty protobuf class.

Recompilation Required?

YES

PBType.outputStreamFromByteRange

Return value type has been changed from `com.google.protobuf.CodedOutputStream` to `org.apache.hbase.thirdparty.com.google.protobuf.CodedOutputStream`.

Package Name

org.apache.hadoop.hbase.types

Effect

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

Reason for change

CDPD-33849

Recommendation

Method signature changed, return type changed to HBase thirdparty protobuf class. If a client application calls this method, it should import the HBase thirdparty protobuf class.

Recompilation Required?

YES

Table.coprocessorService

Added `org.apache.hbase.thirdparty.com.google.protobuf.ServiceException` exception thrown.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: unreported exception `org.apache.hbase.thirdparty.com.google.protobuf.ServiceException` must be caught or declared to be thrown.

Reason for change

CDPD-33849

Recommendation

Method signature changed, thrown Exception type changed to HBase thirdparty protobuf class. If a client application calls this method, and handles this Exception makes sure to import it from the HBase thirdparty protobuf class(es).

Recompilation Required?

YES

ClusterMetrics

Abstract method List `getMasterTasks()` has been added to this interface.

Package Name

org.apache.hadoop.hbase

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getMasterTasks()` in `org.apache.hadoop.hbase.ClusterMetrics`.

Reason for change

HBASE-26730

Recommendation

If a class extends `ClusterMetrics` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

ClusterMetrics

Abstract method List `getUnknownServerNames( )` has been added to this interface.

Package Name

org.apache.hadoop.hbase

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getUnknownServerNames ( )` in `org.apache.hadoop.hbase.ClusterMetrics`.

Reason for change

HBASE-27104

Recommendation

If a class extends `ClusterMetrics` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

ServerMetrics

Abstract method List `getTasks( )` has been added to this interface.

Package Name

org.apache.hadoop.hbase

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getTasks ( )` in `org.apache.hadoop.hbase.ServerMetrics`.

Reason for change

HBASE-26730

Recommendation

If a class extends `ServerMetrics` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `BalanceResponse` balance has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method balance (`BalanceRequest`) in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

HBASE-26147

Recommendation

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `Future<Void> flushAsync ( TableName; List<byte[ ]> )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `flushAsync ( TableName; List<byte[ ]> )` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

HBASE-26867

Recommendation

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `void flushMasterStore ( )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `flushMasterStore ( )` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

HBASE-27028

Recommendation

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method boolean `isReplicationPeerEnabled ( String )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `isReplicationPeerEnabled ( String )` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

HBASE-27448

Recommendation

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method boolean `isReplicationPeerEnabled ( String )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `isReplicationPeerModificationEnabled ( )` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

HBASE-27783

Recommendation

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `List<TableDescriptor> listTableDescriptorsByState ( boolean )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `listTableDescriptorsByState ( boolean )` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

HBASE-27444

Recommendation

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `List<TableName> listTableNamesByState ( boolean )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `listTableNamesByState ( boolean )` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

HBASE-27444

Recommendation

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method `Future<Void> modifyTableAsync ( TableDescriptor; boolean )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `modifyTableAsync ( TableDescriptor; boolean )` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

HBASE-25549

Recommendation

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Admin

Abstract method boolean `replicationPeerModificationSwitch ( boolean; boolean )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `replicationPeerModificationSwitch ( boolean; boolean )` in `org.apache.hadoop.hbase.client.Admin`.

Reason for change

HBASE-27783

Recommendation

If a class extends Admin in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<BalanceResponse>` balance (BalanceRequest) has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `balance ( BalanceRequest )` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

HBASE-26147

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<Void> flush ( TableName; List<byte[ ]> )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `flush ( TableName; List<byte[ ]> )` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

HBASE-26867

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<Void> flushMasterStore ( )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `flushMasterStore ( )` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

HBASE-27028

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<Boolean> isReplicationPeerEnabled ( String )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `isReplicationPeerEnabled ( String )` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

HBASE-27448

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<Boolean> isReplicationPeerModificationEnabled ( )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `isReplicationPeerModificationEnabled ( )` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

HBASE-27783

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<List<TableDescriptor>> listTableDescriptorsByState ( boolean )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `listTableDescriptorsByState ( boolean )` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

HBASE-27444

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<List<TableName>> listTableNamesByState ( boolean )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `listTableNamesByState ( boolean )` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

HBASE-27444

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<Void> modifyTable ( TableDescriptor; boolean )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `modifyTable ( TableDescriptor; boolean )` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

HBASE-25549

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdmin

Abstract method `CompletableFuture<Boolean> replicationPeerModificationSwitch ( boolean; boolean )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `replicationPeerModificationSwitch ( boolean; boolean )` in `org.apache.hadoop.hbase.client.AsyncAdmin`.

Reason for change

HBASE-27783

Recommendation

If a class extends `AsyncAdmin` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncAdminBuilder

Abstract method AsyncAdminBuilder `setRetryPauseForServerOverloaded ( long; TimeUnit )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message:
a client class C is not abstract and does not override abstract method `setRetryPauseForServerOverloaded ( long; TimeUnit )` in `org.apache.hadoop.hbase.client.AsyncAdminBuilder`.

Reason for change

HBASE-26807

Recommendation

If a class extends `AsyncAdminBuilder` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncBufferedMutator

Abstract method int `getMaxMutations ( )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getMaxMutations ( )` in `org.apache.hadoop.hbase.client.AsyncBufferedMutator`.

Reason for change

HBASE-29148

Recommendation

If a class extends `AsyncBufferedMutator` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncBufferedMutatorBuilder

Abstract method `AsyncBufferedMutatorBuilder setMaxMutations ( int )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `setMaxMutations ( int )` in `org.apache.hadoop.hbase.client.AsyncBufferedMutatorBuilder`.

Reason for change

HBASE-29148

Recommendation

If a class extends `AsyncBufferedMutatorBuilder` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncTableBuilder

Abstract method `AsyncTableBuilder<C> setRequestAttribute ( String; byte[ ] )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `setRequestAttribute ( String; byte[ ] )` in `org.apache.hadoop.hbase.client.AsyncTableBuilder<C>`.

Reason for change

HBASE-27657

Recommendation

If a class extends `AsyncTableBuilder` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

AsyncTableBuilder

Abstract method `AsyncTableBuilder<C> setRetryPauseForServerOverloaded ( long; TimeUnit )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `setRetryPauseForServerOverloaded ( long; TimeUnit )` in `org.apache.hadoop.hbase.client.AsyncTableBuilder<C>`.

Reason for change

HBASE-26807

Recommendation

If a class extends `AsyncTableBuilder` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

ColumnFamilyDescriptor

Abstract method `IndexBlockEncoding getIndexBlockEncoding ( )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getIndexBlockEncoding ( )` in `org.apache.hadoop.hbase.client.ColumnFamilyDescriptor`.

Reason for change

HBASE-27314

Recommendation

If a class extends `ColumnFamilyDescriptor` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

ColumnFamilyDescriptor

Abstract method `Compression.Algorithm getMajorCompactionCompressionType ( )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getMajorCompactionCompressionType ( )` in `org.apache.hadoop.hbase.client.ColumnFamilyDescriptor`.

Reason for change

HBASE-25756

Recommendation

If a class extends `ColumnFamilyDescriptor` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

ColumnFamilyDescriptor

Abstract method `Compression.Algorithm getMinorCompactionCompressionType ( )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `getMinorCompactionCompressionType ( )` in `org.apache.hadoop.hbase.client.ColumnFamilyDescriptor`.

Reason for change

HBASE-25756

Recommendation

If a class extends `ColumnFamilyDescriptor` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

ServiceCaller

Abstract method `void call ( S; RpcController; RpcCallback<R> )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method call `( S; RpcController; RpcCallback<R> )` in `org.apache.hadoop.hbase.client.ServiceCaller<S;R>`.

Reason for change

CDPD-33849

Recommendation

Method signature changed, parameter type(s) changed to HBase thirdparty protobuf class(es). If a client application calls this method, it should import the HBase thirdparty protobuf class(es).

Recompilation Required?

YES

TableBuilder

Abstract method `TableBuilder setRequestAttribute ( String; byte[ ] )` has been added to this interface.

Package Name

org.apache.hadoop.hbase.client

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method `setRequestAttribute ( String; byte[ ] )` in `org.apache.hadoop.hbase.client.TableBuilder`.

Reason for change

HBASE-27657

Recommendation

If a class extends `TableBuilder` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Compression.Algorithm

Abstract method `CompressionCodec reload ( Configuration )` has been added to this class.

Package Name

org.apache.hadoop.hbase.io.compress

Effect

Recompilation of a client program may be terminated with the message: a client class C is not abstract and does not override abstract method reload (Configuration) in `org.apache.hadoop.hbase.io.compress.Compression.Algorithm`.

Reason for change

HBASE-26259

Recommendation

If a class extends `Compression.Algorithm` in a client applications, it must implement the abstract method.

Recompilation Required?

YES

Removed APIs in 7.3.2

The following APIs are no longer available for HBase in Cloudera Runtime 7.3.2

Table.batchCoprocessorService

Method removed

Package Name

org.apache.hadoop.hbase.client

Effect

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

Reason for change

CDPD-33849

Recommendation

Method signature changed, return type and parameter types changed to HBase thirdparty protobuf classes. If a client application calls this method, it should import the HBase thirdparty protobuf class(es).

Recompilation Required?

Yes

AuthUtil.getAuthRenewalChore

Method removed

Package Name

org.apache.hadoop.hbase

Effect

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

Reason for change

HBASE-26212

Recommendation

Non-public method. Clients should not depend on non-public methods.

Recompilation Required?

No

Bytes.<init>

Method removed

Package Name

org.apache.hadoop.hbase.util

Effect

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

Reason for change

None

Recommendation

None

Recompilation Required?

No

PBType.encodedLength

Method removed

Package Name

org.apache.hadoop.hbase.types

Effect

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

Reason for change

CDPD-33849

Recommendation

Method signature changed, parameter type changed to HBase thirdparty protobuf class. If a client application calls this method, it should import the HBase thirdparty protobuf class.

Recompilation Required?

Yes

API Compatibility changes in 7.3.2 for Kafka

Removed or Modified APIs in Cloudera Runtime 7.3.2 for Kafka and recommendations for how to handle them.

Apache Version of Kafka in 7.3.1 was 3.4.1 and Apache Version of Kafka in 7.3.2 is 3.9.1

Modified APIs in 7.3.2

The following APIs have been modified for Kafka in Cloudera Runtime 7.3.2

Consumer

Modified

Package Name

org.apache.kafka.clients.consumer

Effect

Recompilation of a client program may be terminated with the message, a client class C is not abstract and does not override abstract method clientInstanceId (Duration) in org.apache.kafka.clients.consumer.

Reason for change

KAFKA-15613

Recommendation

This change affects only the Consumer interface. If you are using the default implementation KafkaConsumer, then no change is required. For custom implementation, see KIP-714: Client metrics and observability.

Recompilation Required?

Yes

Producer

Modified

Package Name

org.apache.kafka.clients.producer

Effect

Recompilation of a client program may be terminated with the message, a client class C is not abstract and does not override abstract method clientInstanceId (Duration) in org.apache.Kafka.clients.producer.Producer.

Reason for change

KAFKA-15613

Recommendation

This change affects only the Producer interface. If you are using the default implementation KafkaProducer, then no change is required. For custom implementation, see KIP-714: Client metrics and observability.

Recompilation Required?

Yes

API Compatibility changes in 7.3.2 for Spark

Removed or Modified APIs in Cloudera Runtime 7.3.2 for Spark and recommendations for how to handle them.

Apache Version of Spark in 7.3.1 was 3.5.4 and Apache Version of Spark in 7.3.2 is 3.5.4

Removed APIs in 7.3.2

The following APIs are no longer available for Spark in Cloudera Runtime 7.3.2

ShuffleExchangeExec.prepareShuffleDependency

Method Removed

Package Name

org.apache.spark.sql.execution.exchange

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

SPARK-51016

Recommendation

None

Recompilation Required?

Yes

API Compatibility changes in 7.3.2 for ZooKeeper

Removed or Modified APIs in Cloudera Runtime 7.3.2 for ZooKeeper and recommendations for how to handle them.

Apache Version of ZooKeeper in 7.3.1 was 3.8.1 and Apache Version of ZooKeeper in 7.3.2 is 3.8.5

Removed APIs in 7.3.2

The following APIs are no longer available for ZooKeeper in Cloudera Runtime 7.3.2

SecurityUtils.createSaslClient

Method Removed

Package Name

org.apache.zookeeper.util

Effect

A client program may be interrupted by NoSuchMethodError exception.

Reason for change

ZOOKEEPER-4889

Recommendation

Client code need to be adjusted to supply ZooKeeper Configuration to the method.

Recompilation Required?

Yes

Unsupported Features

The following section describes the Unsupported Features in the Runtime components.

HDFS

The following HDFS features are currently not supported in Cloudera:

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