

1.0.0

Cloudera Data Warehouse Runtime Release Notes

Date published: 2024-07-26

Date modified: 2024-08-15

CLOUDERA

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

Legal Notice

© Cloudera Inc. 2025. All rights reserved.

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

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

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

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

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

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

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

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

Contents

Cloudera Data Warehouse Runtime changelog..... 4

 Hive changelog for Cloudera Data Warehouse on cloud..... 4

 Impala changelog for Cloudera Data Warehouse on cloud..... 6

 Iceberg changelog for Cloudera Data Warehouse on cloud..... 12

 Hue changelog for Cloudera Data Warehouse on cloud..... 14

Cloudera Data Warehouse on cloud Runtime Changelog

Review the changes introduced in this release of Cloudera Data Warehouse on cloud Runtime. This includes new features, improvements, and fixes for Hive, Impala, Iceberg, and Hue.

Hive changelog for Cloudera Data Warehouse on cloud

Review the changes introduced in Hive for Cloudera Data Warehouse on cloud.

2025.0.19.6 (August 05, 2025)

- No new features or fixes.

2025.0.19.2 (June 04, 2025)

- CDPD-77993: match parquet-hadoop version with downstream

2025.0.19.0 (April 23, 2025) GA

- DWX-19627 HIVE-28620: Query result is cached in case of IOWD if the subquery is not trivial
- DWX-19668 HIVE-28616: OrcReader is not closed in OrcEncodedDataReader
- CDPD-75048 HIVE-13153: SessionID is appended to thread name twice - backport again to fix thread name noise
- CDPD-75556: Call createDmlFragment() when inserting into a partitioned table
- DWX-19504: HIVE-22948: QueryCache: Treat query cache locations as temporary storage
- CDPD-74640: HIVE-25907: IOW Directory queries fails to write data to final path when query result cache is enabled
- CDPD-74640: HIVE-21386: Extend the fetch task enhancement done in HIVE-21279 to make it work with query result cache
- CDPD-74640: HIVE-28530: Set files in thread safe manner in HiveSequenceFileInputFormat
- CDPD-72985: HIVE-27984: Support backward compatibility of hms thrift struct about column stats
- CDPD-71484 HIVE-27884: LLAP: Reuse FileSystem objects from cache across different tasks in the same LLAP daemon
- DWX-18331: HIVE-28121: Use direct SQL for transactional altering table parameter
- CDPD-64754: HIVE-28202: Incorrect projected column size after ORC upgrade to v1.6.7
- DWX-17033: Fixed NPE for UA non BinaryPredicate join condition.
- CDPD-69885: HIVE-28253: Unable to set the value for hplsql.onerror in hplsql mode.
- CDPD-69724: HIVE-28247: Execute immediate `select count(*) from tbl` throwing ClassCastException in hplsql mode.
- CDPD-69648: HIVE-28215: Signalling CONDITION HANDLER is not working in HPLSQL.
- CDPD-69647: HIVE-28214: HPLSQL not using the hive variables passed through beeline using --hivevar option

2024.0.18.4 (December 5, 2024) GA

- DWX-19668 HIVE-28616: OrcReader is not closed in OrcEncodedDataReader
- CDPD-75048 HIVE-13153: SessionID is appended to thread name twice
- CDPD-75556: Call createDmlFragment() when inserting into a partitioned table
- DWX-19504: HIVE-22948: QueryCache: Treat query cache locations as temporary storage
- CDPD-74640: HIVE-25907: IOW Directory queries fails to write data to final path when query result cache is enabled
- CDPD-74640: HIVE-21386: Extend the fetch task enhancement done in HIVE-21279 to make it work with query result cache

- CDPD-74640: HIVE-28530: Set files in thread safe manner in HiveSequenceFileInputFormat
- CDPD-72985: HIVE-27984: Support backward compatibility of hms thrift struct about column stats
- CDPD-71484 HIVE-27884: LLAP: Reuse FileSystem objects from cache across different tasks in the same LLAP daemon

2024.0.18.2-4 (October 4, 2024)

No new features or fixes.

2024.0.18.1-1 (August 15, 2024)

No new features or fixes.

2024.0.18.0-206 (July 26, 2024)

- DWX-18331: HIVE-28121: Use direct SQL for transactional altering table parameter
- CDPD-64754: HIVE-28202: Incorrect projected column size after ORC upgrade to v1.6.7
- DWX-17033: Fixed NPE for UA non BinaryPredicate join condition.
- CDPD-69885: HIVE-28253: Unable to set the value for hplsql.onerror in hplsql mode.
- CDPD-69724: HIVE-28247: Execute immediate `select count(*) from tbl` throwing ClassCastException in hplsql mode.
- CDPD-69648: HIVE-28215: Signalling CONDITION HANDLER is not working in HPLSQL.
- CDPD-69647: HIVE-28214: HPLSQL not using the hive variables passed through beeline using --hivevar option
- CDPD-69170:HIVE-28188:Upgrade Postgres to 42.7.3 to target CVE-2024-1597
- CDPD-7721: HIVE-22489: Reduce Sink operator should order nulls by parameter
- CDPD-68719: upgrade stack version to 7.2.18.0
- DWX-17905: HIVE-28165: HiveSplitGenerator: send splits through filesystem instead of RPC in case of big payload
- DWX-17905: HIVE-24415: HiveSplitGenerator blocks Tez dispatcher
- DWX-17957: HIVE-28167: Full table deletion fails when converting to truncate for Iceberg and ACID tables
- HIVE-27751 Log Query Compilation summary in an accumulated way
- DWX-17928: HIVE-27585: Upgrade kryo serialization lib to latest version (#4570).
- CDPD-67940: HIVE-28111: FSCountersSummary is only printed when LLAP IO is enabled
- CDPD-67793: HIVE-26955: Select query fails when decimal column data type is changed to string/char/varchar in Parquet
- CDPD-65291 :Drop database query failing with Invalid ACL Exception
- CDPD-67350: HIVE-27775: DirectSQL and JDO results are different when fetching partitions by timestamp in DST shift
- CDPD-66862: Upgrade mysql-connector-java in Hive packaging Docker File
- CDPD-65448: HIVE-28006: Materialized view with aggregate function incorrectly shows it allows incremental rebuild
- DWX-17487: HIVE-28051: LLAP: cleanup local folders on startup and periodically
- CDPD-67328: HIVE-28076: Selecting data from a bucketed table with decimal column type throwing NPE.
- CDPD-66902: HIVE-27406: Addendum: Query runtime optimization
- CDPD-61128: HIVE-27490: HPL/SQL says it support default value for parameters but not considering them when no value is passed
- CDPD-62437: HIVE-27692: Explore removing the maintenance task from the embedded Metastore
- CDPD-66930: Recompute scan ranges and stats in ImpalaHdfsScanNode
- CDPD-67205: HIVE-28083: Enable HMS client/query cache for EXPLAIN queries
- CDPD-64442: HIVE-27924: Incremental rebuild goes wrong when inserts and deletes overlap between the source tables
- CDPD-66644: HIVE-28050: Disable Incremental non aggregated materialized view rebuild in presence of delete operations

- CDPD-57121: HIVE-24411: Make ThreadPoolExecutorWithOomHook more awareness of OutOfMemoryError
- CDPD-67063: HIVE-28078: TestTxnDbUtil should generate csv files when we query the metastore database
- DWX-17615: HIVE-28093: Re-execute DAG in case of NoCurrentDAGException
- CDPD-61611: HIVE-27778: Alter table command gives error after computer stats is run with Impala
- CDPD-66267: HIVE-26882: Allow transactional check of Table parameter before altering the Table
- DWX-17274: TestInputFormatReaderDeletes fix
- CDPD-66632: CDPD-50060: HIVE-27114: Include all properties by default
- DWX-16861: CVE fixes for Apache hive JDBC driver
- DWX-17435: Facing dependency issues while using hive JDBC jar with JDBC client other than beeline
- CDPD-66644: HIVE-28050: Disable Incremental non aggregated materialized view rebuild in presence of delete operations
- CDPD-61250: Remove jwtprovider-knox
- CDPD-8806: HIVE-22961: Drop function in Hive should not send request for drop database to Ranger plugin
- CDPD-66527: HIVE-28054: SemanticException for join condition in subquery
- CDPD-63420: HIVE-25183: Inner Join condition parsing error in subquery
- CDPD-66089: HIVE-28000: Hive QL : "not in" clause gives incorrect results when type coercion cannot take place.
- CDPD-66090: HIVE-27161: MetaException when executing CTAS query in Druid storage handler
- DWX-17186 : HIVE-27938: Iceberg: Fix java.lang.ClassCastException during vectorized reads on partition columns
- CDPD-66084: HIVE-27948: Wrong results when using materialized views with non-deterministic/dynamic functions
- CDPD-63299:HIVE-27889: Bump org.apache.santuario:xmlsec from 2.3.0 to 2.3.4
- CDPD-65291 : HIVE-27911 : Drop database query failing with Invalid ACL Exception
- CDPD-65019: HIVE-28009: Shared work optimizer ignores schema merge setting in case of virtual column difference)
- CDPD-61127: HIVE-27489: HPL/SQL does not support table aliases on column names in loops
- CDPD-65805 HIVE-27993 : Netty4 ShuffleHandler: should use one boss thread

Impala changelog for Cloudera Data Warehouse on cloud

Review the changes introduced in Impala for Cloudera Data Warehouse on cloud.

2025.0.19.6 (August 05, 2025)

- IMPALA-14223: Cleanup subdirectories in INSERT OVERWRITE
- IMPALA-14224: Cleanup subdirectories in TRUNCATE
- IMPALA-13778: Update ignored_dir_prefix_list
- CDPD-83059: IMPALA-14074: Warmup metadata cache in catalogd for critical tables
- IMPALA-14130: Remove wait_num_tables arg in start-impala-cluster.py
- IMPALA-13987: Fix stress_catalog_init_delay_ms check in RELEASE
- IMPALA-13850: Wait until CatalogD active before resetting

CDPD-77579: IMPALA-13813: OAuth/JWT Avoid key verification on every rpc call

2025.0.19.2 (June 04, 2025)

- IMPALA-14070: Use checkedMultiply in SortNode.java
- IMPALA-13873: Missing equivalence conjunct in aggregation node with in
- IMPALA-14006: Bound max_instances in CreateInputCollocatedInstances
- IMPALA-13991: Skip CROSS_JOIN rewrite if subquery is in disjunctive
- CDPD-82303: IMPALA-14001: Start EXEC_TIME_LIMIT_S timer after backend execution begins

2025.0.19.0 (April 23, 2025) GA

- IMPALA-13635: Fix wrong expect in test_event_processor_error_global_invalidate
- CDPD-75836: Selectively backport CDPD-75588 to CDWH-2024.0.18.4
- CDPD-45946: IMPALA-13386: Fix test_metastore_service.py after HIVE-27984
- CDPD-72961: IMPALA-13280: Speed up test_under_statement_expression_limit
- IMPALA-13445: Ignore num partition for unpartitioned writes
- IMPALA-13312: Use client address from X-Forwarded-For Header in Ranger Audit Logs
- IMPALA-13310 Add the value of the http 'X-Forwarded-For' header to the runtime profile
- CDPD-73442: IMPALA-13313: Fix ExpireQueries deadlock
- CDPD-72886: IMPALA-13272: Analytic function of collections can lead to crash
- CDPD-72852: IMPALA-13270: Fix IllegalStateException on runtime filter
- CDPD-72923: Build and publish debug docker images
- IMPALA-13137: Add additional client fetch metrics columns to the queries page
- CDPD-66588: IMPALA-12800: Implement hashCode everywhere
- IMPALA-13159: Fix query cancellation caused by statestore failover
- IMPALA-13152: Avoid NaN, infinite, and negative ProcessingCost
- IMPALA-13138: Never smallify existing StringValue objects, only new ones during DeepCopy
- IMPALA-13131: Azure OpenAI API expects 'api-key' instead of 'Authorization' in the request header
- IMPALA-12800: Add cache for isTrueWithNullSlots() evaluation
- IMPALA-13119: Fix cost_initialization at CostingSegment.java
- IMPALA-12616: Fix test_restart_services.py::TestRestart tests for S3
- IMPALA-13129: Move runtime filter skipping at registerRuntimeFilter
- IMPALA-8042: Assign BETWEEN selectivity for discrete-unique column
- IMPALA-12800: Skip $O(n^2)$ ExprSubstitutionMap::verify() for release builds
- IMPALA-13040: (addendum) Inject larger delay for sanitized build
- CDPD-70838: IMPALA-13134: DDL hang with SYNC_DDL enabled when Catalogd is changed to standby status
- DWX-17667: IMPALA-11499: Refactor UrlEncode function to handle special characters
- IMPALA-13040: Add waiting mechanism in UpdateFilterFromRemote
- IMPALA-13018: Block push down of conjuncts with implicit casting on base columns for jdbc tables
- IMPALA-13058: Init first_arrival_time_ and completion_time_ with -1
- CDPD-69477: IMPALA-13054: Avoid revisiting children in QueryStateExpanded
- CDPD-69624: IMPALA-13061: Create query live as external table
- IMPALA-12910: Support running TPCH/TPCDS queries for JDBC tables
- IMPALA-13049: Add dependency management for log4j2 to use 2.18.0
- CDPD-69345: IMPALA-13045: Wait for impala_query_live to exist
- CDPD-68718: IMPALA-12997: Use graceful shutdown for query log tests
- CDPD-69182: IMPALA-13012: Lower default query_log_max_queued
- CDPD-69062: IMPALA-13005: Create Query Live table in HMS
- IMPALA-13024: Ignore slots if using default pool and empty group
- DWX-18112: Temporary fix: Forcibly set CDH-7.2.18.0 dependency
- IMPALA-12657: Improve ProcessingCost of ScanNode and NonGroupingAggregator
- IMPALA-12988: Calculate an unbounded version of CpuAsk
- IMPALA-12980: Translate CpuAsk into admission control slots
- IMPALA-12965: Add debug query option RUNTIME_FILTER_IDS_TO_SKIP
- IMPALA-13004: Fix heap-use-after-free error in ExprTest AiFunctionsTest
- IMPALA-12874: Identify active and standby catalog and statestore in the web debug endpoint
- IMPALA-12564: Prevent Hive loading libfusesupport.so in the minicluster during TSAN runs
- CDPD-69001: IMPALA-12998: Add SHOW_METADATA_TABLES to ignored DDL
- CDPD-68292: IMPALA-12963: Return parent PID when children spawned

- IMPALA-12920: Support ai_generate_text built-in function for OpenAI's chat completion API

2024.0.18.4 (December 5, 2024) GA

- CDPD-75836: Selectively backport CDPD-75588 to CDWH-2024.0.18.4
- CDPD-45946: IMPALA-13386: Fix test_metastore_service.py after HIVE-27984
- CDPD-45946: Fix test_metastore_service.py after HIVE-27984
- CDPD-69857: FENG: Disable test of insert partition with special characters in partition name
- CDPD-72961: IMPALA-13280: Speed up test_under_statement_expression_limit
- CDPD-75562: IMPALA-13469: Deflake test_query_cpu_count_on_insert
- IMPALA-13445: Ignore num partition for unpartitioned writes
- IMPALA-13312: Use client address from X-Forwarded-For Header in Ranger Audit Logs
- IMPALA-13310 Add the value of the http 'X-Forwarded-For' header to the runtime profile
- CDPD-73442: IMPALA-13313: Fix ExpireQueries deadlock

2024.0.18.2-4 (October 4, 2024)

No new features or fixes.

2024.0.18.1-1 (August 15, 2024)

- IMPALA-13270: Addressing IllegalStateException in Complex Views post upgrade
- IMPALA-13272: Stability Improvement for analytic functions on collections

2024.0.18.0-206 (July 26, 2024)

- IMPALA-13137: Add additional client fetch metrics columns to the queries page
- IMPALA-12800: Implement hashCode everywhere
- IMPALA-13159: Fix query cancellation caused by statestore failover
- IMPALA-13152: Avoid NaN, infinite, and negative ProcessingCost
- IMPALA-13138: Never smallify existing StringValue objects, only new ones during DeepCopy
- IMPALA-13131: Azure OpenAI API expects 'api-key' instead of 'Authorization' in the request header
- IMPALA-12800: Add cache for isTrueWithNullSlots() evaluation
- IMPALA-12800: Use HashMap for ExprSubstitutionMap lookups
- IMPALA-13119: Fix cost_ initialization at CostingSegment.java
- IMPALA-12616: Fix test_restart_services.py::TestRestart tests for S3
- IMPALA-13129: Move runtime filter skipping at registerRuntimeFilter
- IMPALA-8042: Assign BETWEEN selectivity for discrete-unique column
- IMPALA-12800: Skip $O(n^2)$ ExprSubstitutionMap::verify() for release builds
- IMPALA-13040: Inject larger delay for sanitized build
- IMPALA-13134: DDL hang with SYNC_DDL enabled when Catalogd is changed to standby status
- IMPALA-11499: Refactor UriEncode function to handle special characters
- IMPALA-13040: Add waiting mechanism in UpdateFilterFromRemote
- IMPALA-13018: Block push down of conjuncts with implicit casting on base columns for jdbc tables
- IMPALA-13058: Init first_arrival_time_ and completion_time_ with -1
- IMPALA-13054: Avoid revisiting children in QueryStateExpanded
- IMPALA-13061: Create query live as external table
- IMPALA-12910: Support running TPCH/TPCDS queries for JDBC tables
- IMPALA-13049: Add dependency management for log4j2 to use 2.18.0
- IMPALA-13045: Wait for impala_query_live to exist
- IMPALA-12997: Use graceful shutdown for query log tests
- IMPALA-13012: Lower default query_log_max_queued
- IMPALA-13005: Create Query Live table in HMS

- IMPALA-13024: Ignore slots if using default pool and empty group
- IMPALA-12657: Improve ProcessingCost of ScanNode and NonGroupingAggregator
- IMPALA-12988: Calculate an unbounded version of CpuAsk
- IMPALA-12980: Translate CpuAsk into admission control slots
- IMPALA-12965: Add debug query option `RUNTIME_FILTER_IDS_TO_SKIP`
- IMPALA-13004: Fix heap-use-after-free error in ExprTest AiFunctionsTest
- IMPALA-12874: Identify active and standby catalog and statestore in the web debug endpoint
- IMPALA-12564: Prevent Hive loading libfessupport.so in the minicluster during TSAN runs
- IMPALA-12998: Add `SHOW_METADATA_TABLES` to ignored DDL
- IMPALA-12963: Return parent PID when children spawned
- IMPALA-12920: Support `ai_generate_text` built-in function for OpenAI's chat completion API
- IMPALA-12969: Release JNI array if DeserializeThriftMsg failed
- IMPALA-12881: Use `getFkPkJoinCardinality` to reduce scan cardinality
- IMPALA-11430: Support custom hash schema for Kudu range tables
- IMPALA-12960: Fix Incorrect RowsPassedThrough Metric in Streaming Aggregation
- IMPALA-12925: Fix decimal data type for external JDBC table
- IMPALA-12899: Temporary workaround for BINARY in complex types
- IMPALA-12602: Unregister queries on idle timeout
- IMPALA-12782: Show info of the event processing in /events webUI
- IMPALA-12852: Make Kudu service start and stop independent
- IMPALA-12426: Switches the duration fields to be stored in decimal seconds.
- IMPALA-12944: Fixes Workload Management Test Flakiness
- IMPALA-12426: Skip Inserting HS2 Operation Queries into the Completed Queries Table
- IMPALA-12626: Add Tables Queried to profile/history
- IMPALA-12264: Add limit on number of HS2 sessions per user.
- IMPALA-12540: Add EventSequence arg to load
- IMPALA-12540: Query Live Table
- IMPALA-12824: Removes the `prettyprint_duration` Built-in Function
- IMPALA-12942: deflake test_virtual_column_file_position_generic
- IMPALA-12487: Skip reloading file metadata for ALTER_TABLE events with trivial changes in StorageDescriptor
- IMPALA-12928: Mask JDBC table property `dbcp.password` for DESC FORMATTED and SHOW CREATE TABLE
- IMPALA-12913: Refactor Workload Management Custom Cluster Tests
- IMPALA-12699: Set timeout for catalog RPCs
- IMPALA-12807: Add support for mold linker
- IMPALA-12915: Use libgtest.so when built with shared libs
- IMPALA-12903: Querying virtual column `FILE__POSITION` for TEXT and JSON tables crashes Impala
- IMPALA-12896: JDBC table must be created as external table
- IMPALA-4545: Simplify test dimension in test_decimal_casting.py
- IMPALA-12898: Tidy up test dimensions of test_scanner.py
- IMPALA-12929: Skip loading HDFS permissions in local-catalog mode
- IMPALA-12930: Fix TestExtDataSources.test_jdbc_data_source failure
- IMPALA-12923: Fix header alignment during horizontal scrolling in query timeline
- IMPALA-12926: Refactor BINARY type handling in the backend
- IMPALA-12819: Sync add-opens targets
- IMPALA-12443: Add catalog timeline for all DDL profiles
- IMPALA-11938: Raised error if NUM_NODES is set to invalid value.
- IMPALA-12802: Support ALTER TABLE for JDBC tables
- IMPALA-12924: Upgrade xmlsec to 2.2.6

- IMPALA-12426: Query History Table
- IMPALA-12883: Support updating the charge on an entry in the cache
- IMPALA-12904: test_type_conversions_hive3 silently passes because of wrongly defined test dimensions
- IMPALA-12914: TestBannedLogMessages no longer requires running cluster
- IMPALA-12818: Intermediate Result Caching plan node framework
- IMPALA-12835: Fix event processing without hms_event_incremental_refresh_transactional_table
- IMPALA-12917: Skip TestEventProcessingError tests requiring tez execution on hive
- IMPALA-12916: Fix test_event_processor_error_global_invalidate test random failure
- IMPALA-12896: Avoid JDBC table to be set as transactional table
- IMPALA-12831: Fix HdfsTable.toMinimalTCatalogObject() failed by concurrent modification
- IMPALA-12832: Implicit invalidate metadata on event failures
- IMPALA-12798: Tolerate floating point precision for regr_intercept() TIMESTAMP test
- IMPALA-12730: Don't use -Weverything for clang-tidy
- IMPALA-12849: Standby catalogd should reject requests from coordinators
- IMPALA-12855: Fix NPE in firing RELOAD events when the partition doesn't exist
- IMPALA-12678: Deflake test_skipping_batching_events
- IMPALA-12878: Report invalid query if finalized
- IMPALA-12860: Invoke validateDataFilesExist for RowDelta operations
- IMPALA-12864: Deflake test_query_log_size_in_bytes.
- IMPALA-12851: Fix AllocWriteIdEvent process issue to add txnId-tableWriteIds mapping
- IMPALA-12815: Support timestamp for scan predicates for external data source table.
- IMPALA-12833: Enabled 'catalogd_ha_reset_metadata_on_failover' by default
- IMPALA-12821: Fix test_create_table_profile_events depend on event-processor
- IMPALA-12825: Install thrift into the impala-python virtualenv
- IMPALA-12834: Add number of concurrent queries to profile
- IMPALA-12859: Ensure query log estimated memory is initialized
- IMPALA-12838: Adds exec_options parameter to single_node_perf_run.py
- IMPALA-12717: Fix error message for missing table descriptor
- IMPALA-12426: QueryStateRecord Refactor
- IMPALA-12827: Fix failures in processing AbortTxnEvent due to aborted write id is cleaned up
- IMPALA-12426: Workload Management Supporting Changes
- IMPALA-12840: Exclude THdfsFileDesc in getJsonCatalogObject
- IMPALA-12793: Create JDBC table without data source
- IMPALA-12823: Fix repeated query not found messages in impalad.INFO logs
- IMPALA-12801: Increase query_log_ default size and bound its memory.
- IMPALA-12426: Remove Usage of Unnecessary this->
- IMPALA-12433: Share buffers among channels in KrpcDataStreamSender
- IMPALA-12824: Adds built-in functions prettyprint_duration and prettyprint_bytes.
- IMPALA-12842: Make node.js installation aware of CPU architecture
- IMPALA-12830: test_webserver_hide_logs_link() could fail in the exhaustive build
- IMPALA-12847: Expose computeScanRangeLocations and computeStats
- IMPALA-11123: Reimplement ORC optimized count star
- IMPALA-9086: Show Hive configurations in /hadoop-varz page
- IMPALA-12796: Add is_footer_only in TFileSplitGeneratorSpec
- IMPALA-12631: Improve count star performance for parquet scans
- IMPALA-12811: Exception during re-analyze can be lost
- IMPALA-12814: Allow use of row_regex for VERIFY_IS_NOT_IN
- IMPALA-11568: Fix SamplingTimeSeriesCounter not correctly cleared
- IMPALA-12816: Remove exclusion of Netty to allow SSL connections to HBase
- IMPALA-12790: Fix overestimation in ScanNode.getInputCardinality

- IMPALA-12803: Fix missing exchange lines in query timeline
- IMPALA-12795: Deflake test_catalog_operation_fields
- IMPALA-12688: Support JSON profile imports in webUI
- IMPALA-12744: Support for regr_count() aggregate/analytic function
- IMPALA-12788: Fix HBaseTable still get loaded even if HBase is down
- IMPALA-12378: Add commons-dbcp2.version replacement
- IMPALA-12578: Pass owner user of database and table to Ranger in GRANT/REVOKE
- IMPALA-12378: Auto Ship JDBC Data Source
- IMPALA-12780: Only show non-default options in the catalog operations page
- IMPALA-12784: Fix bug in ScanNode.getFilteredInputCardinality
- IMPALA-12503: Support date data type for predicates for external data source table
- IMPALA-12783: Nested struct with varlen data crashes
- IMPALA-12726: Simulate large-scale query in TpcdsCpuCostPlannerTest
- IMPALA-12655: Stabilize compute-table-stats.sh
- IMPALA-12770: Fix infinite loop for nested Case expressions
- IMPALA-12781: ARRAY<STRUCT<s: STRING> crashes in top-n
- IMPALA-12718: Provides UTF-8 support for the trim functions
- IMPALA-12778: Fix time strings in catalog operations page
- IMPALA-12642: Fixed unit-test to verify query options for JDBC external table
- IMPALA-12695: Crash with UNION with complex types
- IMPALA-12448: Avoid getting stuck when refreshing a non-existent partition
- IMPALA-12471: skip mysql ext jdbc tests if setup environment fails.
- IMPALA-12763: Union with string struct crashes in ASAN
- IMPALA-12762: Fix cmake error in package building
- IMPALA-12767: Upgrade Guava to 32.0.1 due to CVE-2023-2976
- IMPALA-12463: Batch non-consecutive table events in the event processor
- IMPALA-12745: Skip parallel symbol dumping with RPM/DEB packages
- IMPALA-12125: Support for dumping symbols from RPMs without separate symbols
- IMPALA-12746: Bump jackson.databind to 2.15.3
- IMPALA-12727: Reduce IO threads for non-TARGET_FILESYSTEM filesystems
- IMPALA-12740: Fix TestHdfsJsonScanNodeErrors fails in exhaustive mode
- IMPALA-12711: Fix DDL errors are not shown in impalad logs
- IMPALA-12590: Fix dmesg call during precommit for Ubuntu 20.04
- IMPALA-12584: Enable strict data file access by default
- IMPALA-12038: Switch report_benchmark_results.py to python 3
- IMPALA-12716: Fix timeout thresholds in test_catalog_operations_with_rpc_retry
- IMPALA-12431: Support reading compressed JSON file
- IMPALA-12690: Add TpcdsCpuCostPlannerTest
- IMPALA-12642: Support query options for Impala external JDBC table
- IMPALA-12054: Lazily check Kudu flags in tests
- IMPALA-12582: Fix crash when enabling MIN_MAX RuntimeFilter in Nested Loop Join
- IMPALA-12665: Adjust complete_micro_batch_length to new scratch_batch_->capacity after ScratchTupleBatch::Reset
- IMPALA-12687: Fix key conflicts in tracking in-flight catalog operations
- IMPALA-12356: Fix first ALTER_PARTITION event from Hive could be treated as self event
- IMPALA-12747: Atomic update of execution state
- IMPALA-12719: Reload filemetadata for AlterTable event of type truncate
- IMPALA-12566: Fix RpcMgrKerberizedTest on RedHat 8
- IMPALA-12743: Fix incremental stats are filtered out by HMS due to HIVE-27114

Iceberg changelog for Cloudera Data Warehouse on cloud

Review the changes introduced in Iceberg for Cloudera Data Warehouse on cloud.

2025.0.19.6 (August 05, 2025)

Hive - Iceberg changes

- CDPD-83530: HIVE-28962: Prevent committing outputs if an exception was thrown in the Tez processor.

Impala - Iceberg changes

- IMPALA-14014: Fix COMPUTE STATS with TABLESAMPLE clause
- CDPD-85228/IMPALA-14154: IllegalStateException with Iceberg table with DELETE
- CDPD-87405: IMPALA-14185: Error unnesting nested array from Iceberg with DELETE files

2025.0.19.2 (June 04, 2025)

Hive - Iceberg changes

- DWX-20754: HIVE-28938: Error in LATERAL VIEW with non native tables due to incorrect virtual columns

Impala - Iceberg changes

No new features or fixes.

2025.0.19.0 (April 23, 2025) GA

Hive - Iceberg changes

- CDPD-75084: HIVE-28570: Iceberg: Aggregate queries on translated iceberg table fails post rename.
- DWX-18477: HIVE-28282: Merging into iceberg table fails with copy on write when values clause has a function call.

Impala - Iceberg changes

- CDPD-68442: IMPALA-12970: Fix ConcurrentModificationException for Iceberg table scans
- IMPALA-13432: test_migrated_table_field_id_resolution_complex can fail with different Iceberg version
- CDPD-74017: IMPALA-13364: Schema resolution doesn't work for migrated partitioned Iceberg tables that have complex types
- DWX-17667: IMPALA-11499: Refactor UrlEncode function to handle special characters
- CDPD-69348: IMPALA-13046: Update Iceberg mixed format deletes test
- IMPALA-13002: Iceberg V2 tables with Avro delete files aren't read properly
- IMPALA-12991: Eliminate unnecessary SORT for Iceberg DELETES
- CDPD-68860: IMPALA-13003: Handle Iceberg AlreadyExistsException

2024.0.18.4 (December 5, 2024) GA

- CDPD-68442: IMPALA-12970: Fix ConcurrentModificationException for Iceberg table scans
- IMPALA-13432: test_migrated_table_field_id_resolution_complex can fail with different Iceberg version
- CDPD-74017: IMPALA-13364: Schema resolution doesn't work for migrated partitioned Iceberg tables that have complex types
- CDPD-75084: HIVE-28570: Iceberg: Aggregate queries on translated iceberg table fails post rename.

2024.0.18.2-4 (October 4, 2024)

Hive - Iceberg changes

- DWX-18477: HIVE-28282: (Addendum) Merge query failures with reserved keywords as column names

2024.0.18.1-1 (August 15, 2024)

No new features or fixes.

2024.0.18.0-206 (July 26, 2024)

Hive - Iceberg changes

- DWX-13719: Iceberg positional deletes cache support
- DWX-18477: HIVE-28282: Merging into iceberg table fails with copy on write when values clause has a function call
- CDPD-69704: HIVE-28266: Iceberg: select count(*) from data_files metadata tables gives wrong result
- CDPD-71800: Iceberg: Use 'merge-on-read' as default write mode for v2 tables
- CDPD-70188: HIVE-28278: Iceberg: Stats: IllegalStateException Invalid file: file length 0
- DWX-18213: NFQE SMG Iceberg: Hive select after deletes on table returning inconsistent data
- CDPD-68088: HIVE-28127,HIVE-28203: Exception when rebuilding materialized view with calculated columns on iceberg sources
- CDPD-68087: HIVE-28126: Use added record count in cost model when rebuilding materialized view stored by iceberg
- CDPD-68654: HIVE-28172: Iceberg: Truncate should work for V1 tables with partition evolution.
- CDPD-68285: HIVE-28166: Iceberg: Truncate on branch operates on the main table.
- DWX-17402: HIVE-28069: Iceberg: Implement Merge task functionality for Iceberg tables - Add serde properties for CTAS and CMV
- CDPD-67867: HIVE-28130: Iceberg: Add support for Rename Branch.
- DWX-17957: HIVE-28167: Full table deletion fails when converting to truncate for Iceberg and ACID tables
- DWX-17926: HIVE-28020: Iceberg: Upgrade iceberg version to 1.4.3
- DWX-17926: HIVE-27819: Iceberg: Upgrade iceberg version to 1.4.2
- CDPD-66779: HIVE-28087: Iceberg: Timestamp partition columns with transforms are not correctly sorted during insert
- DWX-17402: HIVE-28069: Iceberg: Implement Merge task functionality for Iceberg tables
- CDPD-66865: HIVE-28119: Iceberg: Allow insert clause with a column list in Merge query not_matched condition
- DWX-16464: HIVE-27653: Iceberg: Add conflictDetectionFilter to validate concurrently added data and delete files
- CDPD-67008: HIVE-28021: Iceberg: Attempting to create a table with a percent symbol fails
- CDPD-67521: HIVE-28015: Iceberg: Add identifier-field-ids support in Hive
- CDPD-67286: HIVE-28102: Iceberg: Invoke validateDataFilesExist for RowDelta operations.
- CDPD-67120: HIVE-28084: Iceberg: COW fix for Merge operation
- DWX-15939: HIVE-27850: Iceberg: Addendum 2: Set runAs user in CompactionInfo
- CDPD-66414: HIVE-28057: Iceberg: Branches with non-lowercase characters can't be accessed.
- DWX-17513: HIVE-27980: Iceberg: Compaction: Add support for OPTIMIZE TABLE syntax
- DWX-17274: Cache Iceberg delete files in LLAP cache
- DWX-15939: HIVE-27850: Iceberg: Addendum: Use fully qualified table name in compaction query
- CDPD-65978: HIVE-27409: Iceberg: table with EXTERNAL type can not use statistics to optimize the query
- CDPD-65863: HIVE-27593: Iceberg: Keep Iceberg properties in sync with HMS

- CDPD-65867: HIVE-27360: Iceberg: Don't create the redundant MANAGED location when creating table without EXTERNAL keyword
- DWX-17270: HIVE-27991: Utilise FanoutWriters when inserting records in an Iceberg table when the records are unsorted
- CDPD-65654: Move the iceberg version dependency to parent pom
- CDPD-65755: HIVE-28016: Iceberg: NULL column values handling in COW mode
- DWX-13719: Iceberg positional deletes cache support

Impala - Iceberg changes

- IMPALA-13046: Update Iceberg mixed format deletes test
- IMPALA-13002: Iceberg V2 tables with Avro delete files aren't read properly
- IMPALA-12991: Eliminate unnecessary SORT for Iceberg DELETES
- IMPALA-13003: Handle Iceberg AlreadyExistsException
- IMPALA-12612: SELECT * queries expand complex type columns from Iceberg metadata tables
- IMPALA-12893: Update tests for Iceberg 1.4.3
- IMPALA-12611: Add support to MAP type Iceberg Metadata table columns
- IMPALA-12600: Schema evolution with equality delete files
- IMPALA-12609: Implement SHOW METADATA TABLES IN statement to list Iceberg Metadata tables
- IMPALA-12809: Iceberg metadata table scanner should always be scheduled to the coordinator
- IMPALA-12894: Fix optimized count(*) for Iceberg tables with dangling delete files
- IMPALA-12729: Allow creating primary keys for Iceberg tables
- IMPALA-12879: Conjunct not referring to table field causes ERROR for Iceberg table
- IMPALA-12893: Specify 'format-version' explicitly in Iceberg tests
- IMPALA-12894: Turn off the count(*) optimisation for V2 Iceberg tables
- IMPALA-12845: Crash with DESCRIBE on a complex type from an Iceberg table
- IMPALA-12610: Support reading ARRAY columns for Iceberg Metadata tables
- IMPALA-12598: Allow multiple equality field id lists for Iceberg tables
- IMPALA-12412: Support partition evolution in OPTIMIZE statement
- IMPALA-12808: test_iceberg_deletes_and_updates failed by timeout error
- IMPALA-12072: Include snapshot id of Iceberg tables in query plan / profile
- IMPALA-12787: Concurrent DELETE and UPDATE operations on Iceberg tables can be problematic
- IMPALA-12764: Fix Iceberg metadata table scan's LIMIT clause
- IMPALA-12742: Add interop tests for DATE-partitioned Iceberg tables
- IMPALA-12765: Balance consecutive partitions better for Iceberg tables
- IMPALA-12742: DELETE/UPDATE Iceberg table partitioned by DATE fails with error
- IMPALA-12708: An UPDATE creates 2 new snapshots in Iceberg tables
- IMPALA-11710: Table properties are not updated in Iceberg metadata files
- IMPALA-12704: Fix NPE when querying empty iceberg table's metadata
- IMPALA-12706: Fix nested struct querying for Iceberg metadata tables
- IMPALA-12495: Describe statement for Iceberg metadata tables

For more information about the core Apache Iceberg 1.4.3 changes, see the [Apache Iceberg Release Notes](#).

Hue changelog for Cloudera Data Warehouse on cloud

Review the changes introduced in Hue for Cloudera Data Warehouse on cloud.

2025.0.19.6 (August 05, 2025)

- DWX-21173: [libsaml] Add custom config for libsaml redirect URL

2025.0.19.2 (June 04, 2025)

No new features or fixes.

2025.0.19.0 (April 23, 2025) GA

- CDPD-72754: vLLM integration with LLama
- DWX-15075: Queries not getting listed in Impala job browser
- CDPD-76594: [Importer] Update Iceberg table syntax for Impala
- CDPD-76409: [ozone] Fix edge cases when volume name start with service_id in path
- CDPD-76282: [Filebrowser] Fix support for reading Parquet files
- CDPD-72496: [fb] Stop file uploads with restrictive file types for old upload handlers
- CDPD-71994: [core] Add the inline documents search back in the documents page
- CDPD-69626: Make semantic search top-k configurable
- CDPD-68357: Multi database support - UI change
- CDPD-61994: [sqlai] Skip loading of sample data for views
- CDPD-58397: Integrate with CML model deployment
- CDPD-57994: [Importer] Fix importer for Parquet, ORC, and Avro formats

2024.0.18.4 (December 5, 2024) GA

- CDPD-58397: Integrate with model deployment
- CDPD-72754: vLLM integration with LLama

2024.0.18.2-4 (October 4, 2024)

No new features or fixes.

2024.0.18.1-1 (August 15, 2024)

No new features or fixes.

2024.0.18.0-206 (July 26, 2024)**New features**

- General availability (GA) of the SQL AI Assistant
- Introduction of task server in Hue and significant improvement in the file upload functionality
- Improvements to the shared Hue service (Preview)

Improvements and bug fixes

- CDPD-68692: [beeswax] fixed the incorrect set_nulls in hive results
- CDPD-68893: [jb] Fix issues with Impala query detailed view in the Job Browser
- CDPD-68891: [jb] Fix issues where Hive was showing up in the job browser even though it was disabled
- CDPD-68683: [tb] Fix an issue where column comments were fetched for individual columns in the Table Browser
- CDPD-67359: [raz][s3] Fix encoding for special chars in directory name listing 1000+ keys
- CDPD-67711: CDPD-65521 DWX-17620 [raz][abfs] Improve url params encoding for RAZ ABFS client
- CDPD-68882: [editor] Replace the editor_autocomplete_timeout setting with disable_source_autocomplete

- CDPD-68875: [editor] Update the coordinator URL pattern in Impala mini job browser
- CDPD-64979: [api] Provide configs to control CORS
- CDPD-67186: [frontend] Advance the 'AND' rule in the definition of 'between' syntax for Hive and Presto
- CDPD-39975: [api] Add new /rename public api
- CDPD-66185: [api] Add new /touch public api
- CDPD-67167: [api] Add new /mkdir public API
- CDPD-65034: Check csrf to see if it has value before setting it up
- CDPD-67174: [jdbc] Add a new Kyuubi jdbc connector
- CDPD-64884: [abfs] Fix none check for ABFSStat attributes
- CDPD-67175: [gs] Add missing gcs-oauth2-boto-plugin package for non-raz scenarios
- CDPD-62397: [core] Default to correct editor choice per user with default_app value
- CDPD-62397: [frontend] Make it possible to configure the default editor to use for all users in Hue
- CDPD-67341: [idbroker] Refactor IDBroker with improved HA support and giving more preference to RAZ when both are configured in Hue