1.5.5

Apache Iceberg Overview

Date published: 2022-07-24
Date modified: 2025-11-08

CLOUDERA

https://docs.cloudera.com/

Legal Notice

Contents

Iceberg introduction.....4
Apache Iceberg Feature Support Matrix.....5

Iceberg introduction

Apache Iceberg is a table format for huge analytics datasets that defines how metadata is stored and data files are organized. Iceberg is also a library that compute engines can use to read/write a table. Iceberg supports concurrent reads and writes on object stores, such as Ozone and on file systems, such as HDFS. HadoopFileIO is supported.

Supported engines and platforms

Querying Apache Iceberg from Apache Hive and Apache Impala is fully supported in Cloudera Data Warehouse on premises, deployed on Cloudera Base on premises version 7.1.9. In this deployment, Cloudera integrates Iceberg with the following engines:

- Apache Flink
- Apache Hive
- Apache NiFi
- Apache Impala
- Apache Spark 3

Querying Iceberg from Hive and Impala is a technical preview in Cloudera Data Warehouse on premises, deployed on Cloudera Base on premises version 7.1.7 or 7.1.8. These versions of Cloudera Base on premises and Cloudera Data Warehouse on premises are not fully interoperable with Iceberg. Iceberg when deployed on these versions is integrated with the following engines only:

- Apache Hive
- Apache Impala

Cloudera supports Iceberg on Red Hat OpenShift and Embedded Container Service (ECS) platforms in Cloudera Data Warehouse on HDFS and Ozone.

Storage

The Hive metastore stores Iceberg metadata, including the location of the table.

Hive metastore plays a lightweight role in the Catalog operations. Iceberg relieves Hive metastore (HMS) pressure by storing partition information in metadata files on the file system/object store instead of within the HMS. This architecture supports rapid scaling without performance hits.

By default, Hive and Impala use the Iceberg HiveCatalog. Cloudera recommends the default HiveCatalog to create an Iceberg table.

Key features overview

You can use Iceberg when a single table contains tens of petabytes of data, and you can read these tables without compromising performance. From Hive and Impala, you can read and write V1 or V2 Iceberg tables. The following features are included:

- Iceberg V2 tables support reading row-level deletes or updates*, making Apache Iceberg ACID-compliant with serializable isolation and an optimistic concurrency model
- Materialized views of Iceberg tables
- Enhanced maintenance features, such as expiring and removing old snapshots
- Performance and scalability enhancements

* The support for delete operations shown in this table is limited to position deletes. Equality deletes are not supported in these releases.

For more information about Iceberg features in Cloudera, see the Iceberg Feature Support Matrix.

Iceberg table security and visualization

Apache Iceberg integrates Apache Ranger for security. You can use Ranger integration with Hive and Impala to apply fine-grained access control to sensitive data in Iceberg tables. Iceberg is also integrated with Data Visualization for creating dashboards and other graphics of your Iceberg data.

Supported ACID transaction properties

Iceberg supports atomic and isolated database transaction properties. Writers work in isolation, not affecting the live table, and perform a metadata swap only when the write is complete, making the changes in one atomic commit.

Iceberg uses snapshots to guarantee isolated reads and writes. You see a consistent version of table data without locking the table. Readers always see a consistent version of the data without the need to lock the table. Writers work in isolation, not affecting the live table, and perform a metadata swap only when the write is complete, making the changes in one atomic commit.

Iceberg partitioning

The Iceberg partitioning technique has performance advantages over conventional partitioning, such as Apache Hive partitioning. Iceberg hidden partitioning is easier to use. Iceberg supports in-place partition evolution; to change a partition, you do not rewrite the entire table to add a new partition column, and queries do not need to be rewritten for the updated table. Iceberg continuously gathers data statistics, which supports additional optimizations, such as partition pruning.

Iceberg uses multiple layers of metadata files to find and prune data. Hive and Impala keep track of data at the folder level and not at the file level, performing file list operations when working with data in a table. Performance problems occur during the execution of multiple list operations. Iceberg keeps track of a complete list of files within a table using a persistent tree structure. Changes to an Iceberg table use an atomic object/file level commit to update the path to a new snapshot file. The snapshot points to the individual data files through manifest files.

The manifest files track data files across partitions, storing partition information and column metrics for each data file. A manifest list is an additional index for pruning entire manifests. File pruning increases efficiency.

Iceberg relieves Hive metastore (HMS) pressure by storing partition information in metadata files on the file system/object store instead of within the HMS. This architecture supports rapid scaling without performance hits.

Related Information

Apache Software Foundation Iceberg Docs
Using Apache Iceberg

Apache Iceberg Feature Support Matrix

Apache Iceberg is a cloud-native, high-performance open table format for organizing petabyte-scale analytic datasets on a file system or object store. Combined with Cloudera, you can build an Open Data Lakehouse architecture for multi-function analytics and deploy large-scale end-to-end pipelines. Cloudera Iceberg supports Hive, Impala, and Spark engines. The Iceberg feature support matrix table lists the Iceberg features and the engines supported.

Open Data Lakehouse on Cloudera simplifies advanced analytics on all data with a unified platform for structured and unstructured data and integrated data services to enable any analytics use case, from ML and BI to stream analytics and real-time analytics. Apache Iceberg is the foundational technology of the Open Lakehouse architecture.

Note: Cloudera supports the v1 and v2 versions of the Iceberg specification for Hive, Impala, and Spark.

The Apache Iceberg format specification describes the following versions of tables:

- v1: Defines large analytic data tables using open format files.
- v2: Specifies ACID-compliant tables, including row-level deletes and updates.

The following tabs list the latest Iceberg features, supported engines, and their earliest supported version of Cloudera and Cloudera Data Services:

For Features

Table 1: Iceberg feature support matrix

Iceberg feature	Hive	Impala	Spark
Create table	✓	✓	✓
Select	✓	✓	✓
Insert	✓	✓	✓
Update	✓	✓	✓
Delete	✓	✓	✓
Merge	✓	✓	✓
Schema evolution	✓	✓	✓
Partition evolution	✓ (Hidden partition)	✓ (Hidden partition)	✓
Drop partition	✓	✓	✗
Transaction support	✓	✓	✓
Parquet support	✓	✓	✓
Compaction	✓	✓	✓
Equality deletes	✓ (read-only)	✓ (read-only)	✓ (read-only)
Positional deletes	✓	✓	✓ (read and write)
Merge-on-read	✓	✓	✓
Copy-on-write	✓	✗	✓

Iceberg feature	Hive	Impala	Spark
Optimistic concurrency control	✓	✓	✓
Time travel	✓	✓	✓
Rollback	✓	✓	✓
Optimized Row Columnar (ORC)	✓	✓ (read-only)	✓
Avro support	✓	✓ (read-only)	✓
Snapshot isolation level	✓	✓	✓
Serializable isolation level	✗	✗	✓
Branching	✓	✗	✓
Tagging	✓	✗	✓
Fine-grained access control (Ranger)	✓ (HWC)	✓	✗
Table statistics	✓	✓	✓
Partition statistics	✗	✓ (no Number of Distinct Values)	✗
Geospatial types	✓	✗	✗
Partition pruning	✓	✓	✓
Z-Ordering	✗	✓	✓
Bucket map join	✓	✗	✓ (Storage partitioned join)
Migration from external table	✓	✓	✓
Load data inpath	✓	✓	✗

For Hive

Table 2: Hive feature availability

Iceberg feature	Form factor	Version introduced in
Create	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Select	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Insert	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Update	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Delete	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Merge	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Schema evolution	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Partition evolution	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Drop Partition	Cloudera on cloud	7.3.1
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.8.5
	Cloudera Data Warehouse on premises	1.5.4
Transaction support	Cloudera on cloud	7.2.18

Iceberg feature	Form factor	Version introduced in
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Parquet support	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Compaction	Cloudera on cloud	7.3.1 500
	Cloudera Base on premises	7.3.1 500
	Cloudera Data Warehouse on cloud	1.9.1
	Cloudera Data Warehouse on premises	1.5.4 SP1
Equality deletes	Cloudera on cloud	7.3.1 500
	Cloudera Base on premises	7.3.1 500
	Cloudera Data Warehouse on cloud	1.9.1
	Cloudera Data Warehouse on premises	1.5.4 SP1
Positional deletes	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Merge-on-read	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Copy-on-write	Cloudera on cloud	7.3.1 500
	Cloudera Base on premises	7.3.1 500
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Optimistic concurrency control	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Time travel	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Rollback	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4

Iceberg feature	Form factor	Version introduced in
Optimized row columnar support	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Avro support	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.1
	Cloudera Data Warehouse on premises	1.5.4
Snapshot isolation level	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Avro support	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.1
	Cloudera Data Warehouse on premises	1.5.4
Branching	Cloudera on cloud	7.3.1 500
	Cloudera Base on premises	7.3.1 500
	Cloudera Data Warehouse on cloud	1.9.1
	Cloudera Data Warehouse on premises	1.5.4
Tagging	Cloudera on cloud	7.3.1 500
	Cloudera Base on premises	7.3.1 500
	Cloudera Data Warehouse on cloud	1.9.1
	Cloudera Data Warehouse on premises	1.5.4
Fine-grained access control (Ranger)	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Table statistics	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Geospatial types	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Partition pruning	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3

Iceberg feature	Form factor	Version introduced in
	Cloudera Data Warehouse on premises	1.5.4
Bucket map join	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Migration from external table	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4
Load data inpath	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.9.3
	Cloudera Data Warehouse on premises	1.5.4

For Impala

Table 3: Impala feature availability

Iceberg feature	Form factor	Version introduced in
Create	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Select	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Insert	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Update	Cloudera on cloud	7.3.1.500
	Cloudera Base on premises	7.3.1.500
	Cloudera Data Warehouse on cloud	1.8.5
	Cloudera Data Warehouse on premises	1.5.4
Delete	Cloudera on cloud	7.2.18
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.8.1
	Cloudera Data Warehouse on premises	1.5.4
Merge	Cloudera on cloud	7.3.1.500
	Cloudera Base on premises	7.3.1.500

Iceberg feature	Form factor	Version introduced in
	Cloudera Data Warehouse on cloud	1.10.1
	Cloudera Data Warehouse on premises	1.5.5
Schema evolution	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Partition evolution	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Drop Partition	Cloudera on cloud	7.3.1.500
	Cloudera Base on premises	7.3.1.500
	Cloudera Data Warehouse on cloud	1.8.5
	Cloudera Data Warehouse on premises	1.5.4
Single statement transaction support	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Parquet support	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Major compaction	Cloudera on cloud	7.2.18
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.8.5
	Cloudera Data Warehouse on premises	1.5.4
Sophisticated compaction	Cloudera on cloud	7.3.1.500
	Cloudera Base on premises	7.3.1.500
	Cloudera Data Warehouse on cloud	1.10.1
	Cloudera Data Warehouse on premises	1.5.5
Equality deletes	Cloudera on cloud	7.3.1.500
	Cloudera Base on premises	7.3.1.500
	Cloudera Data Warehouse on cloud	1.9.1
	Cloudera Data Warehouse on premises	1.5.4
Positional deletes	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Merge-on-read	Cloudera on cloud	7.2.17

Iceberg feature	Form factor	Version introduced in
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Optimistic concurrency control	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouseon premises	1.5.2
Time travel	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Rollback	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Optimized row columnar support	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Avro support	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Snapshot isolation level	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Fine-grained access control (Ranger)	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Table statistics	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Partition statistics	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2

Iceberg feature	Form factor	Version introduced in
Partition pruning	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Z-Ordering	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2
Migration from external table	Cloudera on cloud	7.2.18
	Cloudera Base on premises	7.3.1
	Cloudera Data Warehouse on cloud	1.7.1
	Cloudera Data Warehouse on premises	1.5.2
Load data inpath	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Warehouse on cloud	1.6.1
	Cloudera Data Warehouse on premises	1.5.2

For Spark

Table 4: Spark feature availability

Iceberg feature	Form factor	Version introduced in
Create	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Select	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Insert	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Update	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Delete	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19

Iceberg feature	Form factor	Version introduced in
	Cloudera Data Engineering on premises	1.5.2
Merge	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Schema evolution	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Partition evolution	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Transaction support	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Parquet support	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Compaction	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Equality delete (read-only)	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Positional deletes	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Merge-on-read	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Copy-on-write	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9

Iceberg feature	Form factor	Version introduced in
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Optimistic concurrency control	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Time travel	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Rollback	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Optimized row columnar support	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Avro support	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Snapshot isolation level	Cloudera on cloud	7.3.1
	Cloudera Base on premises	7.3.1
	Cloudera Data Engineering on cloud	1.22
	Cloudera Data Engineering on premises	1.5.4-h5
Serializable isolation level	Cloudera on cloud	7.3.1
	Cloudera Base on premises	7.3.1
	Cloudera Data Engineering on cloud	1.22
	Cloudera Data Engineering on premises	1.5.4-h5
Branching	Cloudera on cloud	7.3.1
	Cloudera Base on premises	7.3.1
	Cloudera Data Engineering on cloud	1.22
	Cloudera Data Engineering on premises	1.5.4-h5
Tagging	Cloudera on cloud	7.3.1
	Cloudera Base on premises	7.3.1
	Cloudera Data Engineering on cloud	1.22
	Cloudera Data Engineering on premises	1.5.4-h5
Table statistics	Cloudera on cloud	7.2.17

Iceberg feature	Form factor	Version introduced in
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Partition pruning	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2
Z-Ordering	Cloudera on cloud	7.3.1
	Cloudera Base on premises	7.3.1
	Cloudera Data Engineering on cloud	1.22
	Cloudera Data Engineering on premises	1.5.4-h5
Bucket map join	Cloudera on cloud	7.2.18
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.20
	Cloudera Data Engineering on premises	1.5.2
Migration from external table	Cloudera on cloud	7.2.17
	Cloudera Base on premises	7.1.9
	Cloudera Data Engineering on cloud	1.19
	Cloudera Data Engineering on premises	1.5.2

Related Information

Cloudera Data Warehouse Runtime Release Notes