

Impala workload management in Data Warehouse Public Cloud (Preview)

Date published: 2024-07-26

Date modified: 2024-07-30

Legal Notice

© Cludera Inc. 2024. All rights reserved.

The documentation is and contains Cludera 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 Cludera software may be found within the documentation accompanying each component in a particular release.

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

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

Cludera, Cludera Altus, HUE, Impala, Cludera Impala, and other Cludera 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.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cludera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Contents

Legal Notice	2
Contents	3
About Impala workload management	4
Table format	4
Use cases	9
How Impala workload management works	10
Impala coordinator startup flags	10
Known Issue	12

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

About Impala workload management

Cloudera Data Warehouse (CDW) Public Cloud provides you the option to enable logging Impala queries on an existing Virtual Warehouse or while creating a new Impala Virtual Warehouse. By logging the Impala queries in CDW, you gain increased observability of the workloads running on Impala, which you can use to improve the performance of your Impala Virtual Warehouses.

Note: *This feature is in technical preview and not recommended for use in production deployments. Cloudera recommends that you try this feature in test and development environments.*

This feature significantly enhances the query profiling capabilities. You can have Impala archive crucial data from each query's profile into dedicated database tables, namely the query history table and live query table. These tables are part of the `sys` database and are designed to store valuable information that can later be queried using any Impala client, providing a consolidated view of reports from previously executed queries.

The query history table, `sys.impala_query_log`, proves particularly useful when dissecting workloads for in-depth analysis of query performance. Unlike the limitations associated with query profiles, which are only available to the client that initiated the query, the query history table offers a comprehensive solution for querying completed queries without the need to parse the text of each query profile. Additionally, the query history table provides a comprehensive view across all Impala coordinators.

The Impala query information is stored indefinitely in the `sys.impala_query_log` table whereas the `sys.impala_query_live` table reflects the in-memory state of all Impala coordinators. Actively running and recently completed queries are stored in the `sys.impala_query_live` table. Data is removed from this table once it is persisted in the `sys.impala_query_log` table or if the coordinator is restarted. Therefore, it is possible that some of the records are momentarily duplicated in both these tables.

Because the `sys.impala_query_live` table is stored only in memory, recently completed queries that still need to be persisted to the `sys.impala_query_log` table are lost if the coordinator crashes. However, if the coordinator shuts down gracefully, then the recently completed queries are stored in the `sys.impala_query_log` table and are not lost.

Table format

The Impala query history and live query system tables contain the following columns:

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

CLUDERA TECHNICAL PREVIEW DOCUMENTATION

Column Name	Description	Data Type	Sample Value
cluster_id	String specified through the Impala startup flag to uniquely identify an instance	string	cluster-123
query_id	Impala assigned query identifier	string	214d08bef0831e7a:3c65392400000000
session_id	Impala assigned session identifier	string	ea4f661af43993d8:587839553a41adb8
session_type	Client session type	string	HIVESERVER2
hiveserver2_protocol_version	Version of the HiveServer (HS2) protocol that was used by the client when connecting	string	HIVE_CLI_SERVICE_PROTOCOL_V6
db_user	Effective user on the cluster	string	csso_name
db_user_connection	Username from an authenticated client	string	csso_name
db_name	Name of the database being queried	string	default
impala_coordinator	Name of the coordinator for the query	string	coord-22899:27000
query_status	Status of the query when it completes	string	OK
query_state	Final state of the query	string	FINISHED
impala_query_end_state	Final Impala state of the query	string	FINISHED
query_type	Type of the query	string	QUERY
network_address	Client IP address and port	string	127.0.0.1:40120

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

CLUDERA TECHNICAL PREVIEW DOCUMENTATION

start_time_utc	Time when the query started. Time zone is in UTC	timestamp	2024-07-17 17:13:46.414316000
total_time_ms	Difference between the query end time and start time, in milliseconds	decimal(18,3)	136.121
query_opts_config	List of query options stored as a single string containing comma-separated values of key-value pairs	string	TIMEZONE=America/Los_Angeles,CLIENT_IDENTIFIER=Impala Shell v4.4.0a1 (04bdb4d) built on Mon Nov 20 10:49:35 PST 2023
resource_pool	Name of the resource pool for the query	string	default-pool
per_host_mem_estimate	Size, in bytes of the per-host memory estimate	bigint	5
dedicated_coord_mem_estimate	Size, in bytes of the dedicated coordinator memory estimate.	bigint	4
per_host_fragment_instances	Comma-separated string listing each host and its fragment instances	string	myhost-1:27000=1,m yhost-2:27001=2
backends_count	Count of the number of backends used by this query	integer	2
admission_result	Result of the admission (not applicable to DDLs)	string	Admitted immediately
cluster_memory_admitted	Cluster memory, in bytes that was admitted	integer	4
executor_group	Name of the executor group	string	executor_group

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

CLUDERA TECHNICAL PREVIEW DOCUMENTATION

executor_groups	List of all executor groups including the groups that were considered and rejected as part of Workload Aware Auto Scaling	string	executor_group1, executor_group2...
exec_summary	Full text of the executor summary	string	
num_rows_fetched	Number of rows fetched by the query	bigint	6001215
row_materialization_rows_per_sec	Count of the number of rows materialized per second	bigint	3780
row_materialization_time_ms	Time spent materializing rows converted to milliseconds	decimal(18,3)	1.58
compressed_bytes_spilled	Count of bytes that were written (or spilled) to scratch disk space	bigint	241515
event_planning_finished	Event from the timeline. The value represents the number of milliseconds since the query was received.	decimal(18,3)	27.253
event_submit_for_admission	Event from the timeline. The value represents the number of milliseconds since the query was received.	decimal(18,3)	30.204
event_completed_admission	Event from the timeline. The value represents the number of	decimal(18,3)	30.986

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

CLUDERA TECHNICAL PREVIEW DOCUMENTATION

	milliseconds since the query was received.		
event_all_backends_started	Event from the timeline. The value represents the number of milliseconds since the query was received.	decimal(18,3)	31.969
event_rows_available	Event from the timeline. The value represents the number of milliseconds since the query was received.	decimal(18,3)	31.969
event_first_row_fetched	Event from the timeline. The value represents the number of milliseconds since the query was received.	decimal(18,3)	135.175
event_last_row_fetched	Event from the timeline. The value represents the number of milliseconds since the query was received.	decimal(18,3)	135.181
event_unregister_query	Event from the timeline. The value represents the number of milliseconds since the query was received.	decimal(18,3)	141.435
read_io_wait_total_ms	Total read I/O wait time converted to milliseconds	bigint	15.091
read_io_wait_mean_	Average read I/O wait	bigint	35.515

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

ms	time across executors converted to milliseconds		
bytes_read_cache_total	Total bytes read from the data cache	bigint	45823
bytes_read_total	Total bytes read	bigint	745227
pernode_peak_mem_min	Minimum value of all the per-node peak memory usages	bigint	5552846
pernode_peak_mem_max	Maximum value of all the per-node peak memory usages	bigint	5552846
pernode_peak_mem_mean	Mean value of all the per-node peak memory usages	bigint	5552846
sql	SQL statement as provided by the user	string	SELECT db_user, total_time_ms from impala_query_log where query_state = 'EXCEPTION';
plan	Full text of the query plan	string	
tables_queried	Comma-separated string containing all the tables queried in the SQL statement	string	

Use cases

A consolidated view of reports from previously executed queries can be useful in the following use cases:

- For collecting the history of all queries run and reported by a user or by date and time
Example: `SELECT db_user, start_time, end_time, total_time_ms, sql from impala_query_log order by db_user;`
- For collecting the frequently-run queries
Example: `SELECT db_user, query_status, query_state, sql from impala_query_log where query_state = 'EXCEPTION';`
- For reporting queries that are running over 10 minutes

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

Example: `SELECT db_user, total_time_ms, sql from impala_query_log where total_time_ms > 600000;`

How Impala workload management works

To use this feature, enable Impala query logging while creating a new Virtual Warehouse or by editing an existing one by selecting the **Log Impala Queries (Preview)** option. By default, the **Log Impala Queries (Preview)** option is off.

You can then configure the Impala coordinator using specific startup flags to store query history. Impala manages the table serving as a centralized repository for all query histories across databases. Completed queries are periodically inserted into this table based on a preconfigured interval.

This feature streamlines the process of query history management, providing a more accessible and comprehensive way to analyze and retrieve information about completed queries

Note: This feature is available only on CDW running Impala image versions 2024.0.18.0 and higher.

Important: The following query types are not written into the query logging tables:

- SET
- SHOW
- USE
- DESCRIBE

Impala coordinator startup flags

Each Impala coordinator runs an SQL statement on startup to create the query logging tables. The following table lists the Impala startup coordinator flags that you can configure:

Name	Data Type	Default	Description
cluster_id	string		Specifies an identifier string that uniquely represents this cluster. This identifier is included in both tables and is used as a table partition for the <code>sys.impala_query_log</code> table.
query_log_shutdown_deadline_s	number (seconds)	30	Hidden flag. Number of seconds to wait for the queue of completed queries to be carried into the query history table before timing out and continuing the shutdown

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

CLUDERA TECHNICAL PREVIEW DOCUMENTATION

			<p>process.</p> <p>The query history table drain process runs after the shutdown process completes, therefore the max shutdown time is extended by the value specified in this flag.</p>
workload_mgmt_user	string	impala	Specifies the user that will be used to insert records into the query history table.
query_log_write_interval_s	number (seconds)	300	Number of seconds to wait before inserting completed queries into the query history table. Allows for batching inserts to help avoid small files.
query_log_max_queued	number	5000	<p>Maximum number of completed queries that can be queued before they are written to the query history table.</p> <p>This flag operates independently of the 'query_log_write_interval_m' flag. If the number of queued records reaches this value, the records will be written to the query log table no matter how much time has passed since the last write.</p> <p>A value of 0 indicates no maximum number of queued records.</p>
query_log_max_sql_length	number	16777216	<p>Maximum length of an SQL statement that will be recorded in the completed queries table.</p> <p>If an SQL statement with a length longer than this specified value is executed, the SQL inserted into the completed queries table will be trimmed to this length.</p> <p>Any characters that require escaping will have their backslash character counted towards this limit.</p>
query_log_max_plan_length	number	16777216	Maximum length of the SQL query plan that will be recorded in the completed queries table. If a query plan has a length longer than this value, the plan inserted into the completed queries table will be trimmed to this length.

This document has been released as part of a technical preview for features described herein. Technical preview components are provided as a convenience to our customers for their evaluation and trial usage. These components are provided 'as is' without warranty or support. Further, Cloudera assumes no liability for the usage of technical preview components, which should be used by customers at their own risk.

			Any characters that require escaping will have their backslash character counted towards this limit.
query_log_request_pool	string		Specifies a pool or queue used by the queries that insert into the query log table. Empty value causes no pool to be set.

Known Issue

[DWX-18554](#): You cannot access the `sys.impala_query_live` table when Impala cannot schedule the query.

Workaround: None.