

Trino performance tuning

Date published:

Date modified:

CLOUDERA

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

Performance tuning for Trino Virtual Warehouse..... 4
Increasing the number of preferred nodes for caching files..... 4
Configuring resource management..... 4
Configuring spill memory..... 5

Performance tuning for Trino Virtual Warehouse

This section provides guidance on performance tuning for Trino Virtual Warehouses, covering topics such as increasing preferred nodes for caching files, configuring resource management to address memory issues, and setting up spill memory properties for memory-intensive queries.



Note: Trino is available as a Technical Preview offering starting from the Cloudera Data Warehouse on premises 1.5.5 SP1 release. It is available as a GA offering in the Cloudera Data Warehouse on premises 1.5.5 SP2 release.

Increasing the number of preferred nodes for caching files

Query processing in Trino is distributed into different stages where tasks are processed by different nodes in the cluster. You can limit the number of hosts that are required to process these tasks with the `fs.cache.preferred-hosts-count` property.

If you are using the Hive or Iceberg connector and if the workload from these tables are high, then a high throughput is required, which can be achieved by increasing the value of the `fs.cache.preferred-hosts-count` connector property in the connector configuration details.

The value must be increased as concurrency increases. The default value of 2 may not be sufficient if you are running a high workload. For example, if you are running 40 nodes with a concurrency of 5, consider increasing the value of `fs.cache.preferred-hosts-count` to 20. You can increase this value up to a maximum limit of 40 (maximum worker nodes allowed for caching files).

Note that the value depends on the workload that you are trying to query through the connector and on the number of parallel queries that you expect to run simultaneously. Therefore, tune the parameter according to your workload needs.

Configuring resource management

If you are experiencing memory issues and notice queries failing, it is recommended that you increase the allocated query memory for the Trino Coordinator and Worker.

Procedure

1. Log in to the Cloudera Data Warehouse service as DWAdmin.
2. Go to the **Virtual Warehouses** tab, locate the Trino Virtual Warehouse and click  Details . The **Trino Virtual Warehouse Details** page is displayed.
3. Click the **Configurations** tab, select Trino Coordinator, and select `trino-config-properties` from the Configuration files drop-down menu.
4. Search for the following properties and modify the values as recommended:
 - `query.max-memory` — Maximum amount of user memory that can be used by a query across the entire cluster.
Recommended value: `query.max-memory-per-node` * number of worker nodes. For example, for 40 nodes, the value is 2560 GB.
 - `query.max-memory-per-node` — Maximum amount of user memory that can be used by a query on a worker node.

Recommended value: 64 GB

KEY	VALUE
query.max-memory	2560
query.max-memory-per-node	64

5. Click Apply Changes.
6. From the **Configuration** tab, click Trino Worker and then select trino-config-properties from the **Configuration files** drop-down menu.
7. Search for the same memory related properties that are specified in Step 4 and set the appropriate values.
8. Click Apply Changes.

Configuring spill memory

During memory intensive queries, if you notice issues related to the spill memory, you can configure certain properties to facilitate running of queries that require memory exceeding per query or per node limits.

Procedure

1. Log in to the Cloudera Data Warehouse service as DWAdmin.
2. Go to the **Virtual Warehouses** tab, locate the Trino Virtual Warehouse and click Details .
The **Trino Virtual Warehouse Details** page is displayed.
3. Click the **Configurations** tab, select Trino Worker, and select trino-config-properties from the Configuration files drop-down menu.
4. Click Add Custom Configuration and add the following properties along with the recommended values:
 - query-max-spill-per-node — Maximum spill space that can be used by a single query on a single worker node.
Recommended value: 100 GB
 - max-spill-per-node — Maximum spill space that can be used by all queries on a single worker node.
Recommended value: 275 GB

KEY	VALUE
max-spill-per-node	275
query-max-spill-per-node	100



Note: The spill space is shared with the Filesystem cache (fs-cache).

5. Click Apply Changes.