

Magic Comments for Presto Queries

Magic comments are instructions inserted into Presto SQL queries and given by the database driver to override the database optimization that occurs with standard Presto queries. The database driver follows certain rules for adding the magic comments to the database query and the magic comments override any execution plan the query optimizer might select for a query.

Use magic comments in cases when you know more about the database content and structure than the query optimizer. It can be more efficient to use magic comments to improve execution time.

- [Magic Comment Parsing Decimal Literals as Double](#)
- [Magic Comment Distributed Sort](#)
- [Magic Comment Time Partitioning Range](#)
- [Magic Comment Limiting Execution Time](#)
- [Magic Comment Extend Query Runtime Beyond Configured Limits](#)
- [Magic Comment to improve performance for result output](#)

Magic Comment Parsing Decimal Literals as Double

A Presto session-level property that controls whether the decimal literals(ex. 1.2) are parsed as decimal or double.

The `parse_decimal_literals_as_double` session property is supported as a magic comment.

Decimal literals without an explicit type specifier (e.g. 1.2) are treated as the values of the DOUBLE type by default.

This magic comment can disable the behavior and use decimal literals as the values of the DECIMAL type. For example:

```
-- set session parse_decimal_literals_as_double = 'false'  
SELECT * FROM large_table, small_table  
WHERE small_table.id = large_table.id
```

Magic Comment Distributed Sort

Presto session property that when enabled allows the sort operator to execute in parallel on multiple nodes in the cluster.

The `distributed_sort` session property is supported as a magic comment.

Distributed sort allows sorting of data, which exceeds `query.max-memory-per-node`. Distributed sort is enabled through the `distributed_sort` session property, or `distributed-sort` configuration property set in `etc/config.properties` of the coordinator. Distributed sort is enabled by default.

When distributed sort is enabled, the sort operator executes in parallel on multiple nodes in the cluster. Partially sorted data from each worker node is then streamed to a single worker node for a final merge. This technique allows to utilization of memory of multiple worker nodes for sorting. The primary purpose of the distributed sort is to allow for the sorting of data sets that don't normally fit into single node memory. Performance improvement can be expected, but it won't scale linearly with the number of nodes, since the data needs to be merged by a single node.

```
-- set session distributed_sort = 'true'  
SELECT * FROM large_table, small_table  
WHERE small_table.id = large_table.id
```

Magic Comment Time Partitioning Range

A TD-specific setting to customize partition size for INSERT/CTAS queries.

The `time_partitioning_range` session property is supported as a magic comment.

```
-- set session time_partitioning_range = 'value'
```

The value can be:

- none

- no time partitioning
- number with unit
 - allowed units are h, d, mo, q, y (hour, day, month, quarter, year)
 - ex) 12h, 2d, 1mo, 1q, 1y

Magic Comment Limiting Execution Time

If you want to override any other limits for a single query, you can apply the following query hint at the beginning of the query:

```
-- set session query_max_execution_time = '12h'
SELECT
  COUNT(*)
FROM
  ALL_EVENTS
;
```

If limits are set in your account, and you write a query with a hint, then the smallest limit between account, resource-group, query-hint, and global is used. Supported time measures are d (days), h (hours), m (minutes), s (seconds).

If the limit specified in the query hint is longer than the configured limit, then the shorter configured limit still applies. For example, if a 1-hour limit is configured for ad hoc queries, then you can use the magic comment to limit a query to 5 minutes (5m) but not to increase the limit to 2 hours (2h).



All Presto queries are limited to 6 hours runtime, regardless of any account-level policy or query hint.

Magic Comment Extend Query Runtime Beyond Configured Limits

You can run a query in a special *overcommit* resource pool if you want to extend the runtime beyond the query's normal limits. In this case, the limit specified in the query hint is greater than the query's normal limit.

For example, if there is a 30-minute limit configured on scheduled queries, you can run a longer query with the following query hints:

```
-- set property resource_group = 'overcommit'
-- set session query_max_execution_time = '1d'
```

Magic Comment to improve performance for result output

To improve the job performance of queries that produce massive results, *result_output_redirect* is used to enhance the performance of Presto jobs. By default, *result_output_redirect* is set to 'true'.

```
-- set session result_output_redirect='true'
```

This impacts performance even though the query results are ready; TD Presto is waiting for the worker node to complete its sequential operations. To mitigate this, Treasure Data now uses *result_output_direct*, which redirects the query result to S3 in parallel, thereby improving the performance of the queries.

LIMITATION: There are SQL queries where Hints do not improve performance, including:

- When a query includes ORDER BY
- When a query is CREATE TABLE or INSERT TABLE
- When a query includes uppercase letters
- When a query has duplicated column names
- When a query has NULL without explicit type