

Loading

The load settings on the datamodel allow you to set all detailed options on the Celonis Engine loads.

Cache retention

Save changes

You can configure after which time entries should be compressed and how long entries should be kept in the hot-cache when they are not used. Higher usually means better performance but higher memory usage. Note that you have to reload the data model for the setting to become active.

Compress entries after

30

minutes (Default: 30)

Keep entries in cache for

60

minutes (Default: 60)

Number of dimensions from which caching is used. Lower values increase the performance but also increase the memory consumption.

4

Columns

Cache eviction

Save changes

You can configure thresholds to trigger the removal of entries from the hot-cache. Thresholds represent usage percentage of the system memory. Upon the reach of higher threshold, entries are removed until usage of system memory is reduced to the lower threshold. Note that you have to reload the data model for the setting to become active.

Stop removing entries when memory usage is

(Default: 0.5)

Start removing entries when memory usage is

(Default: 0.7)

1. The cache retention defines how long the results are kept in hot-cache when they are not used.

2. The cache eviction defines thresholds when to start and stop removing entries from the hot-cache.

Enable query caching

Save changes

By enable query caching, long running queries whose execution time exceeds one second, will be executed and cached while loading the actual data model. Multiple executions of a query are accelerated as well. Note that you have to reload the data model for the setting to become active. Please note, the query cache allocates memory in the according heap of the application server.

☒ Enable data model query caching

Query cache size

150

MB

Max. cache warm up time

15

Minutes

Load scheduling

Save schedule

Schedule for loading the complete data model.

☐ Activate automatic re-loading

Partial Load scheduling

Save partial load schedule

Schedule for loading partially the data model. Which tables are part of the partial load can be configured in the table editor. The following tables are currently part of a partial load:

_CEL_P2P_ACTIVITIES

☐ Activate automatic partial re-loading

3. Enabling Query Caching will cache queries running longer than one second. Size and warm up time can be configured

4. Load scheduling allows you to define continuous engine reloads. You can choose pre-set time intervals for the reloads, create a custom CRON job or trigger the reload from a database table.

5. Partial load scheduling allows you to define continuous partial reloads of specified tables.

Enable precaching

Save changes

By enable precaching, long running operators whose execution time exceeds two seconds, will be precached while loading the actual data model. Note that you have to reload the data model for the setting to become active.

☒ Enable data model precaching

Distributed calculation

Save changes

Select a server for the execution of this data model.

Log level

Save changes

Select a log level for the execution of this data model. Note that you have to reload the data model for the setting to become active.

DEBUG

6. The Celonis Engine smartly tracks long running operators and uses them to restore the cache during a data model reload. This is called precaching. You can deactivate precaching to save resources.

7. Distributed calculation specifies which Compute Server the data model should be loaded on in case of a Multi-Server-Deployment

8. The Log level specifies the granularity of the Compute Logs

Limit variants for clustering

Save changes

The CLUSTER_VARIANTS and the ESTIMATE_CLUSTER_PARAMS operators perform very expensive computations and require excessive memory and CPU resources. To avoid out of memory errors and unresponsive analyses, these operators can be limited to certain number of distinct variants.

Maximum number of variants for clustering:

100000 (default)

variants

Load column

Optimize load of table

Optimized load allows you to speed up the load of data model significantly.

9. Settings to limit the CLUSTER_VARIANTS and ESTIMATE_CLUSTER_PARAMS operators

10. Load column allows you to exclude columns in the data tables. Those columns will not be loaded to the Celonis Engine. This can significantly reduce load times.