

Configuring Aggregate Life Cycle Settings

You can configure aggregate life cycle settings to control the number of system-defined aggregates (both demand-defined and prediction-defined).

Procedure

To configure these settings:

1. At the top of the Design Center, choose **Settings**.
2. On the left, choose **Aggregates**.
3. Configure the aggregate life cycle settings (described below).
4. Save your changes.



Setting a value of **0** for the **Retention Limit** and **Excess Amount** properties is unsupported and will revert both settings to their default values. A setting of **1** for both properties is valid.

Concurrent Queries Per Organization

The maximum number of system aggregate queries that can run simultaneously. The AtScale engine runs these queries when building or rebuilding instances of aggregate tables. If this limit is exceeded, aggregate-table instances are not built until prior builds are complete.

For example, suppose you leave the value at the default of 1 and there are three aggregate tables (A, B, and C) defined for a cube. When the batch build process for the cube runs, the AtScale engine submits a query to gather the data for the instance of aggregate table A. Only after the instance of A is built will the engine then submit a query for building the instance of aggregate table B, and so on.

The default value is: 1

Retention Limit Per Cube

AtScale will begin cleaning up less-used system-defined aggregate tables when this number is exceeded in a cube.

The default value is: 100

Excess Amount Per Cube

Number of per Cube, system-defined aggregates that are allowed in excess of the retention limit per cube. Aggregates are pruned back to the target number when this threshold is reached.

The default value is: 10

Remove Unused Aggregates

Disable aggregate definitions with zero utilizations after this amount of time (in days). Applied when the total aggregate count is less than the RETENTION LIMIT per cube.

The default value is: 45

Cache Settings

You can also use the following aggregate life cycle settings:

- ▲ `aggregates.dataWarehouseCacheTableRequests.enabled`: Toggles whether to cache data warehouse table requests. Default value is true.



Note: This setting applies only for data warehouses that use the Spark engine; for details, see the [data warehouse support matrix](#).

- ▲ `aggregates.dataWarehouseCacheTableRequests.maximumRowCount`: The maximum number of rows for a data warehouse table request to be cacheable, if caching is enabled. Default value is 50000.

They can be set in the following ways:

- ▲ For an organization, go to Settings, choose Engine in Organization Settings, and set them as needed in Custom Settings. For details, see [Changing Advanced Engine Settings](#).
- ▲ For a cube, you can set them as described in [Cube Level Configuration Settings](#).