

Other Settings For Both System-Defined And User-Defined Aggregates

A number of features of system-defined and user-defined aggregates are enabled or affected by single settings.



Your user ID must be a super user.

To access these settings, go to **Settings** > **Engine**. Use your browser Find to search for each setting.

AGGREGATES.CREATE.BUILDFROMEXISTING

Set to **True** to allow new aggregate tables to be built from data that is in an existing aggregate table. This option does not affect subsequent builds.

The AtScale engine continuously assesses the quality of the aggregate-table definitions that it has generated. If it determines that a new definition is needed, by default the first instance of that definition is built from a query against raw data, even if that definition is based on a current aggregate-table definition.

Use this setting to allow the first instance of a new definition to be built from the data that is already in an instance of another definition. Allowing the first instance to be built in this way speeds up the build process.

For example, suppose that the engine decides to supersede the aggregate-table definition AggDef1 by creating the new definition AggDef2, which is based on AggDef1. If this setting is set to True, the build of the first instance of AggDef2 will include data from the current instance of AggDef1. If the instance requires data that is not in the current instance of AggDef1, the engine queries raw data to gather it.

Non-incremental aggregates tables can be built only from non-incremental aggregate tables, while incremental aggregate tables can be built only from incremental aggregate tables.

The default value is True.

AGGREGATE.CREATE.SECURITYDIMENSIONS.ENABLED

Controls whether or not aggregates should be built when an attribute is part of a security dimension. Default value is false.

AGGREGATES.CREATION.TIMEOUT

Specify the maximum length of time to allow per DDL statement that the engine uses to create an aggregate instance. Aggregates that are refreshed with full builds require one DDL statement. Aggregates that are refreshed with incremental builds require one DDL statement per partition.

AGGREGATES.DROP.PURGE.ENABLED

This setting should only be enabled in specific circumstances. For more information, view this article in the Knowledge Base: <https://customers.atscale.com/s/article/How-to-add-the-purge-directive-to-an-aggregate-drop-table-statement>

AGGREGATES.DROP.TIMEOUT

Specify the maximum length of time to allow per DDL statement that the engine uses to drop an aggregate instance. The default value is 1.minute; restart is required. The setting is valid for SQL data sources. BigQuery ignores it, and uses its default 50 seconds timeout.

AGGREGATE.INCREMENTALUPDATES.IMMUTABLE.ENABLED

Set to **True** to enable incremental builds and rebuilds of aggregates that use joins on rarely changing dimensions. For more about such incremental builds, see [About Incremental Rebuilds](#).

AGGREGATES.ORC.COMPRESS

Specify which compression method to use. This setting is applicable only if you set the value of `AGGREGATES.TABLECONFIG.PREFERREDSTORAGEFORMAT` to "orc".

AGGREGATES.TABLECACHE.ENABLED

 To use this feature with AtScale, you must have a special AtScale license.

Enable the in-memory aggregate caching feature. The default value is false.

AGGREGATES.TABLECACHE.EXTRASUPPORTEDFUNCTIONS.ENABLED

Allow the use of the in-memory aggregate cache when the query contains functions from the set of extra functions (for example, LIKE). The default value is true. Can be used when the `AGGREGATES.TABLECACHE.ENABLED` setting (see above) is set to 'true'.

AGGREGATES.TABLECONFIG.PREFERREDSTORAGEFORMAT

Specify the storage format for data in aggregate tables, if you have a preference. Possible values: orc, parquet, rcfile, textfile, none. Specify "none" to allow the engine to decide which format to use.

TABLES.CREATE.PARTITIONS.ENABLED

Set to **True** to enable the AtScale engine to partition table types that can be partitioned. Default = `True`.



Partitioned aggregates are not supported when the data warehouse that you are using is an instance of Google BiqQuery.

If the value of this setting is not `True`, then the `AGGREGATES.CREATE.PARTITION.USERDEFINEDAGGREGATE.ENABLED` and `AGGREGATES.CREATE.PARTITION.SYSTEMDEFINEDAGGREGATE.ENABLED` engine settings have no effect, even when they are set to `True`. For details, see [Engine Settings for User-Defined Aggregates Only](#) and [Settings that Control Partitioning of System-Defined Aggregate Tables](#), respectively.