

Engine Settings For User-Defined Aggregates Only

Use these settings to control partitioning of user-defined aggregates.



Note: AtScale 2021.2.0 and later supports Google BigQuery Aggregate Table partitioning by Integer. AtScale 7.0.0 - 2021.1.0 support Google BigQuery Aggregate table partitioning by only Date and DateTime columns.

To access these settings:

1. Make sure you are logged in as super user.
2. Choose **Settings** from the main navigation, and go to **Engine**.
3. Use your browser Find to search for each setting.

AGGREGATES.CREATE.PARTITION.USERDEFINEDAGGREGATE.ENABLED

Set to **True** to enable the AtScale engine to partition user-defined aggregates. For this setting to have an effect, the setting `TABLE.CREATE.PARTITIONS.ENABLED` must be set to True. Default = **True**

TABLE.CREATE.PARTITIONS.MAXIMUMESTIMATEDNUMBEROFPARTITIONS

Specifies the maximum estimated number of partitions to consider acceptable for an instance of a user-defined aggregate. The creation of an instance will fail if the estimate is larger or cannot be calculated. Default = 800

TABLE.CREATE.PARTITIONS.MAXIMUMDYNAMICPARTITIONSFACTOR

Specifies the maximum number of partitions per table for underlying data warehouse engines. Some engines have their own maximum number of partitions allowed per table. Before being passed to engines, this value is multiplied by `AGGREGATE.CREATE.PARTITIONS.MAXIMUMESTIMATEDNUMBEROFPARTITIONS`. Default = 1.2

TABLE.CREATE.PARTITIONS.MAXIMUMDYNAMICPARTITIONSPERNODEFACTOR

Specifies the maximum number of partitions per node for underlying data warehouse engines. Some engines have a maximum number of table partitions per node. Before being passed to engines, this value is multiplied by `AGGREGATE.CREATE.PARTITIONS.MAXIMUMESTIMATEDNUMBEROFPARTITIONS`. Default = 0.2

More Information

[Partitioned User-Defined Aggregates](#)