

Settings That Control Partitioning Of System-Defined Aggregate Tables

One of these settings enables the AtScale engine to partition instances of aggregate tables that it defines. The other setting lets you set a threshold for determining when to partition them.



In AtScale 7.0 and later, AtScale supports aggregate table partitioning for Google BigQuery using columns of type Date or DateTime for partition columns. In AtScale 2021.2.0 AtScale supports partitioning Google BigQuery aggregate tables by Integer columns.

AGGREGATES.CREATE.PARTITION.SYSTEMDEFINEDAGGREGATE.ENABLED

Set to `True` to enable the AtScale engine to partition system-defined aggregates. For this setting to have an effect, the setting `TABLES.CREATE.PARTITIONS.ENABLED` must be set to `True`.

AGGREGATES.CREATE.PARTITION.SYSTEMDEFINEDAGGREGATE.THRESHOLD

Specify minimum number of rows per partition. The AtScale engine divides the estimated cardinality of a proposed system-defined aggregate table by the estimated number of partitions. If the estimated number of rows per partition does not meet or exceed this threshold, the engine will not partition the aggregate table. This value prevents the engine from creating too many partitions per aggregate table, as query processing times can increase if the number of partitions becomes too high. This value also prevents the engine from creating not enough partitions per aggregate table, as a small number of very large partitions can also cause query processing times to increase. In both cases, the advantages of partitioning are negated. Default = 50000.0.

More Information

- ▲ The two settings above can also be applied on cube level. For details, see [Cube Level Configuration Settings](#).
- ▲ [Model Additional Levels in a Hierarchy](#)
- ▲ [Edit a Level](#)