

About Hinted Aggregate Tables

Hinted aggregates are a type of user-defined aggregate table; you determine whether to allow the AtScale engine to create such aggregate tables on query data sets. You also control the definition of a hinted aggregate with the SQL statement that defines the corresponding query data set. The join keys, measures, and degenerate dimensions that the SQL statement references will be included in the hinted aggregate table for the query data set. If you change the references in that SQL statement, the definition of the hinted aggregate changes and is put into effect when the instance of the aggregate table is rebuilt.

To allow the AtScale engine to create a hinted aggregate table for a query data set, edit an existing query data set. Select **Allow aggregates**, then select **Hint to the engine to materialize this query data set in an aggregate**. After the project is republished, the AtScale engine will build the hinted aggregate table.



This feature is only available with query data sets on the main Design Center Cube Canvas. It is not available with query data sets on dimension canvases.

When the AtScale engine creates a hinted aggregate table on a query data set, some queries that are run against the query data set can run more efficiently against the hinted aggregate table instead. This can save some queries from requiring a fully materialized query data set.

This feature is not a form of a materialized view that you may be familiar with from a traditional RDBMS; AtScale will not write the query data set results to a table. Instead, AtScale writes an aggregate table using the join keys, degenerate dimensions, and measures derived from the query data set. The resulting aggregate table can be much smaller than the results of the query data set.

For example, take the following query data set:

```
SELECT Prod_SKU, Color, Sum(Sales) FROM Sales GROUP BY 1, 2
```

In this example, where the model contains a color dimension and Sales measure, enabling the hinted aggregate feature will create an aggregate consisting of Color and Sum(Sales), which contain significantly fewer rows than the source query data set results contain.

Hinted aggregate tables can be built incrementally. As when you allow incremental aggregate tables to be built for fact data sets, you specify an incremental indicator and a grace period. See [About Incremental Rebuilds](#) for details.

As you can for other types of aggregate table, you can monitor the use of hinted aggregates, the success or failure of builds of hinted aggregates, and other information about them from the **Aggregate Definitions** and **Aggregate Instances** pages. To access those pages, click **Aggregates** at the top of the Design Center.



Unhandled `NULL` values in key columns will result in incomplete aggregate tables and unexpected query results. See the [Custom Empty Member](#) feature for more details.

Partitioned Hinted Aggregates

AtScale applies partitioning when the hinted aggregate includes key columns that match keys on the cube's partitioning keys list. The hinted aggregate will have partitioning applied as per the Cube level settings. This behavior can be disabled by setting the configuration option `aggregates.create.partition.hintedAggregate.enabled` to `false`. The default value for this setting is `true`. It can also be set on cube level as described in [Cube Level Configuration Settings](#).

For more information, see [Partitioned System-Defined Aggregate Tables](#)