

When The AtScale Engine Creates Demand-Defined Aggregates

Demand-defined aggregates are a type of system-defined aggregate that the AtScale engine creates according to various criteria, such as the history of queries against a model and statistics about the use of existing demand-defined aggregates.

Suppose that a client application, such as Tableau or Microsoft Excel, issues a query against a model that was designed in and deployed from AtScale.

The AtScale engine parses the query, noting the metrics and dimensions that are included in it. The engine can then find out whether there is an aggregate table that will satisfy the query. There are two possible outcomes of this search: an appropriate aggregate table does or does not exist.

When An Appropriate Aggregate Table Does Not Exist

When a query requests aggregated data and no aggregate table yet exists that can satisfy the query, the AtScale engine takes four actions:

1. Runs the query against the data.
2. Adds the query to a log in which it compiles a history of the queries against the model. The engine uses this history to refine its estimates of query patterns, an estimate that it uses when deciding which aggregate tables to define or redefine as it continuously optimizes query performance by creating new instances.
3. Determines the best aggregate-table definition or definitions to create to satisfy both the current query and similar queries in the future. The engine does this by examining the metrics and dimensions in the query, the history of queries against the same model, and statistics for the model. A candidate definition could therefore contain metrics that the current query does not reference.

For example:

Suppose that this query requested the count of items (metric) that were sold per customer gender (dimensional attribute) for a particular product line (dimensional attribute). The AtScale engine would include that metric and, if the cost of doing so were negligible, it might also include additional metrics, such as the total dollar sales amount of the sales and the number of sales transactions.

With very little up-front cost, a new aggregate table based on this definition could satisfy not only the current query but possible future queries against these other metrics.

The AtScale engine does not include metrics from more than one fact dataset when creating a definition. If a query references metrics that are in two fact datasets, the AtScale engine considers whether to create two definitions: one for the metrics in one dataset, and another for the metrics in the other dataset.

4. Determines whether the cardinality of an instance of an aggregate table definition would be less than or equal to the limit set by the compression factor. This factor specifies the ratio of the cardinality of the fact dataset to the cardinality of the instance. By default, the ratio is 3 to 1, which means that the AtScale engine will not define an aggregate table for which an instance would have a cardinality greater than 1/3 of the cardinality of the corresponding fact dataset.

For example:

If a fact dataset has one billion rows, an aggregate-table instance created from the proposed definition would have to contain no more than 333,333,333 rows. If the instance would have more rows than that, the AtScale engine would not define the aggregate table.

When An Appropriate Aggregate Table Does Exist

When a query requests aggregated data and an aggregate table that can satisfy the query does exist, the AtScale engine does the following:

- ▲ Examines the query in the context of the query history for the model.
- ▲ Examines statistics for the model.

Using the information that it gathers, the engine decides to take one of the following actions:

- ▲ Create an additional aggregate table definition.
- ▲ Supersede the existing definition that is used to satisfy the query. If the engine decides to supersede the existing definition with a new one, it could choose between narrowing the definition (removing metrics, perhaps because those metrics are not being used) or widening it (adding metrics that the query history shows are also used from the fact dataset and aggregated by the same dimensional attributes).
- ▲ Take both of the preceding actions, creating additional aggregate-table definitions and superseding existing definitions.

A Note About Degenerate Dimensions

The AtScale engine can use aggregate instances that contain complete degenerate dimensions to populate filters in BI tools like Tableau and Microsoft Excel. If an aggregate instance is complete for a degenerate dimension (meaning the instance contains all of the values for it), then querying the aggregate can be faster than querying the fact dataset.

For example, suppose there is an aggregate instance that includes the complete degenerate dimension Color and the metric Order Quantity. In Tableau, adding the Color dimension to the Filters card opens the Filter dialog, which displays the colors in the dimension. Behind the scenes, the AtScale engine can query the aggregate to retrieve the values for populating the Filter dialog.

Figure 1. The Filter dialog in Tableau, displaying all of the values for the degenerate dimension Color

The screenshot shows the Tableau interface with a data table and a 'Filter [Color]' dialog box. The data table has the following content:

Order Quar..	Order Month	
1	January	7,929
	February	7,994
	March	8,145
2	April	8,768
	May	9,594
	June	9,564
3	July	6,346
	August	6,699
	September	6,617
4	October	7,169
	November	7,152
	December	9,190

The 'Filter [Color]' dialog box is open to the 'General' tab. It shows the following options:

- Select from list
- Custom value list
- Use all

The list of values to filter is:

- Black
- Blue
- Multi
- Red
- Silver
- White
- Yellow

Buttons: All, None, Exclude

Summary:

- Field: [Color]
- Selection: Selected 0 of 8 values
- Wildcard: All
- Condition: None
- Limit: None

Buttons: Reset, Apply, Cancel, OK

However, if an aggregate instance contains any non-degenerate dimensions, then the AtScale engine will not use it to populate filters.

More Information

[SQL Hints to Control the Use and Generation of Aggregate Tables](#)