Managing Aggregates

This section explains how AtScale creates and uses aggregate tables (or aggregates for short) to optimize query performance, how you can define your own aggregates, and how administrators can monitor and configure aggregates.

- About Aggregates AtScale incorporates the data-warehousing concept of aggregate tables. Such tables most often contain measures from one or more fact datasets and include aggregated values for these measures. (There are dimension-only aggregate tables.) The aggregation of the data is at the level of one or more dimensional attributes or, if no dimensional attributes are included, the aggregated data is a total of the values for the included measures.
- Changing the Schema for Aggregate Tables AtScale creates its aggregate tables in your data warehouse in a schema that a system administrator creates before deploying AtScale. It is possible to change the schema after AtScale is deployed.
- Engine Settings for Aggregates AtScale provides a number of Engine settings that enable you to configure the creation of aggregates.
- Excluding Specific Attributes From Aggregates You can exclude specific attributes from being used as grouping keys in aggregations. These settings can be applied to both level attributes and secondary dimensional attributes.
- Disabling the Creation of System-Defined Aggregates for a Dataset For some fact datasets, you might want to disable the creation of system-defined aggregates entirely.
- Defining Aggregates Yourself You can define your own aggregates for use cases that fall outside of those covered by system-defined aggregates.
- Monitoring Aggregate Usage You can view what aggregates have been created, when they were last updated, and how often they have been used by a query. You can also see the status of instances and a history of instances for each aggregate.
- Rebuilding Aggregates Manually You can manually perform an initial build or a rebuild all aggregates for a published model.
- Setting Properties to Allow Incremental Rebuilds of Aggregates If you want to allow the AtScale engine the option of performing incremental rebuilds of the aggregates for a model, you must change settings in the properties of the fact dataset and dimensions of the model.
- Performing Full Rebuilds of Incremental Aggregates Incremental aggregates process only new windows of fact data rows, rather than processing all of the data. In some cases, you may want to reprocess all of the data to ensure that aggregates are accurate. For example, if the data of a dimension dataset has changed significantly or if you have older data that missed a processing window.
- Handling NULL Values to Prevent Incomplete Aggregate Tables and Unexpected Query Results
- Inbound SQL Hints to Control the Use and Generation of Aggregate Tables
- Inbound SQL Session Parameters