

Types Of Aggregate Tables In AtScale

There are two types of aggregate table that the AtScale engine can use: System-Defined aggregate tables and User-Defined aggregate tables. System-Defined tables are further divided into Demand-Defined and Prediction-Defined. User-Defined tables are further divided into manual and hinted.

System-Defined Aggregates

The AtScale engine is able to create two different types of aggregates, **Demand-Defined** and **Prediction-Defined**.

Demand-Defined

This type of aggregate is created by the AtScale engine according to various criteria, such as the history of queries against a cube and statistics about the use of existing Demand-Defined aggregates.

For more about this type of Demand-Defined aggregates, see [When the AtScale Engine Creates Demand-Defined Aggregates](#).

Prediction-Defined

The AtScale engine creates these aggregates because, after examining cubes immediately after those cubes are published, it anticipates that the aggregates will be useful for some queries. The purpose of creating these aggregates is to improve the performance of queries issued by client BI software faster than is possible with the creation of Demand-Defined aggregates alone.

There are three different varieties of **Prediction-Defined** aggregates and they are Defined by the AtScale engine when a project is published. When a project is republished the engine revises the definitions of these aggregates, according to the changes that were made in the corresponding cubes.

1. An **all-member** aggregate table is Defined per fact dataset in a cube, if you enable the AtScale engine to define aggregates of this type.

It contains all of the measures from the fact dataset, and contains none of the related dimensions.

2. A **dimension-only** aggregate is specifically useful for filters in BI client software, such as Tableau and Microsoft Excel.

The compression ratio for a proposed dimension-only aggregate, defined as the number of rows in the full dimension dataset divided by the number of rows in the proposed aggregate for a level in the dimensional hierarchy, must be equal to or greater than a ratio that you can configure. You must enable the AtScale engine to define aggregates of this type.

3. A **super aggregate** is defined per fact dataset in a cube, if you allow the AtScale engine to define aggregates of this type.

It contains all keys in a fact dataset. It also contains all degenerate dimensions for which a fact dataset contains values at all levels in the dimensional hierarchies.

- ▲ This type of Prediction-Defined aggregate is useful during the rebuilding of aggregate instances for a cube, as they can be used for rebuilding other aggregates and for avoiding the need to run queries to rebuild instances against fact tables in the connected data warehouse.
- ▲ The AtScale engine can also use super aggregates to build new Demand-Defined aggregates.

For example, if a query runs against columns A, B, and C, and a super aggregate contains columns A, B, C, D, and E, then a Demand-Defined aggregate based on the query can be built directly from data in the super aggregate, again avoiding a query directly against the fact table to populate the Demand-Defined aggregate.

For more information, see [Settings for Prediction-Defined Aggregates](#).

User-Defined Aggregates

There are two different types of User-Defined aggregate tables:

- ▲ **Manual:** These are aggregate tables that you define for certain use cases that System-Defined aggregates do not cover. For more about User-Defined aggregates, see [When to Define Your Own Aggregate Tables](#)
- ▲ **Hinted:** These are a type of aggregate table that you can allow the AtScale engine to define on a query data set. For more information, see [About Hinted Aggregate Tables](#).

More Information

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