

When To Define Your Own Aggregate Tables

AtScale recommends that you rely on the AtScale engine to generate aggregate tables for you. There are use cases that fall outside of those covered by system-defined aggregates. For these cases you can create user-defined aggregates (UDA).

When To Create UDA

If you require aggregate tables that contain any of the types of dimensional attributes or measures listed below, you must define the aggregate tables manually:

- ▲ Measures on dimensions

The AtScale engine does not generate aggregate tables for measures that are local to a dimension only (known as a 'secondary metrical attribute' in the cube model).

- ▲ Non-additive measures

The AtScale engine does not generate aggregate tables for non-additive measures, which are useful for distinct counts. The reason is that such an aggregate table defined for one query would not be usable by other queries. Here is an example:

- ▲ Suppose that some applications run a query (Query 1) against your data to look up distinct counts of users (a non-additive measure) aggregated by month, and other applications run a query (Query 2) that looks up distinct counts of users by week.
- ▲ The engine would not generate an aggregate table for Query 1 because no other query on the distinct counts of users would be able to use that table. For the same reason, the engine would not generate an aggregate table for Query 2. You would need to create aggregate tables manually.

How Aggregates Are Selected

For UDAs to be selected by the AtScale Query Planner, the following criteria must be present:

1. The UDA must be eligible by containing all of the required model attributes required by the query planner. This includes Dimension Name and Key columns used in the inbound query's SELECT clause, WHERE clause, and GROUP BY clause.
2. UDA's using dimensional attributes from multiple dimensions must specify at least one "simple measure" (for example, Sum, Count, Min, or Max) from the fact tables that connect the dimensions. Note the following:
 - ▲ A "simple measure" uses one of the following aggregation functions: SUM, COUNT, MIN, or MAX.
 - ▲ UDA's containing only a DISTINCT COUNT or AVERAGE Measure to connect related fact datasets will not be used by In-bound queries.

3. The UDA must be the best alternative to other tables (aggregates, raw fact, raw dimension) as ranked by up-to-date statistics collected by AtScale. For example, if the selected base table has fewer rows than the competing UDAs, the query planning will select the base table over the UDA.

More Information

- ▲ UDAs are created in the Design Center. For details, see [Defining Aggregates Yourself](#).
- ▲ After you define an aggregate table, the AtScale engine creates the first instance of the aggregate table when you publish the corresponding project.
- ▲ To check that the instance has been created after you published the project, go to the Aggregates page in the Design Center. There you can find a list of definitions, a history of their instances, and also information about build failures (if problems occur when creating the aggregate instance). For details, see [Monitoring Aggregate Usage](#).