About Join Behavior

When a query involves data from two different datasets, AtScale uses the relationships defined in the cube to plan how to join the underlying tables in the connected data warehouse. This section explains the join behavior considered by the AtScale query engine.

> **Note:** The actual SQL join type chosen by the AtScale engine is based on the optimal path to the data requested in the query. When this documentation talks about SQL join types, it is only to illustrate the matching and filtering behavior associated with the relationships defined in your AtScale cubes.

- ▲ About Right vs Left Join Behavior The default behavior is to filter out the fact data rows that do not have a matching entry in the dimension. In SQL, this is referred to as right join behavior, meaning the dimension dataset is used as the basis of comparison when combining data from multiple tables in the connected data warehouse.
- ▲ About Inner vs Outer Join Behavior The default behavior is to exclude dimension members that do not have corresponding fact rows. In SQL, this is often referred to as inner join behavior, meaning there must be a match in both tables for a result to be included.