

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.

For example, if this Sales fact dataset was being joined to this Customer dimension, then only customers A, B, C, and D would be in the result (customer E and F would be excluded because those customers do not have records in the fact data).

| Sales (Fact) | | | | Customer (Dimension) | | |
|--------------|---------|--------|------------|----------------------|--------|-----|
| TransID | Date | Amount | CustomerID | CustomerID | Gender | Age |
| 1 | 9/24/15 | 350 | A | A | F | 22 |
| 2 | 9/24/15 | 25 | B | B | M | 33 |
| 3 | 9/24/15 | 160 | C | C | F | 40 |
| 4 | 9/25/15 | 210 | D | D | F | 19 |
| | | | | E | M | 28 |
| | | | | F | M | 46 |

In some cases, you may want to see all dimension members regardless if they have fact data or not. You can configure per dimension hierarchy if you want this outer join type behavior instead. In this case, customer E and F would be in the result, but the sales amount would be NULL (empty) for those customers.