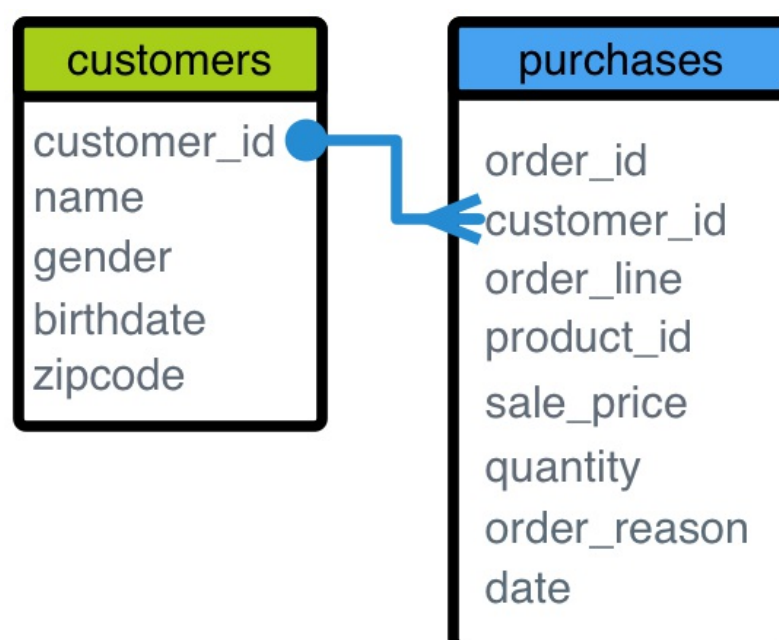


One-To-Many Relationships

When modeling data in a star schema format, dimension-to-fact relationships are typically one-to-many. This means that each record in the fact dataset can link to one (and only one) record in the dimension dataset, but a record in the dimension dataset can be associated with many fact records.

For example, a purchase is associated with exactly one customer, but a customer can have many purchases.



When you model your dimensions and relationships in AtScale, the default assumption is that your source data and tables adhere to the star schema model. The dimension datasets have a primary key (one or more columns that can uniquely identify a dimension record), and the fact dataset has a corresponding foreign key that can be used to combine the two tables at query time.

If your data does adhere to the typical star schema model, then the only thing you need to do when modeling your dimensions is to make sure that the key level of your dimension hierarchies are marked as **Unique**. This helps AtScale optimize queries that involve one-to-many joins.

In the one-to-many star schema model, if the dimension has duplicate keys, then the query results would only include the first match for a given fact record. The duplicates would be discarded from the final results. For example, if there were duplicate customer_id records in the example above, then a given purchase record would only match to one of them.