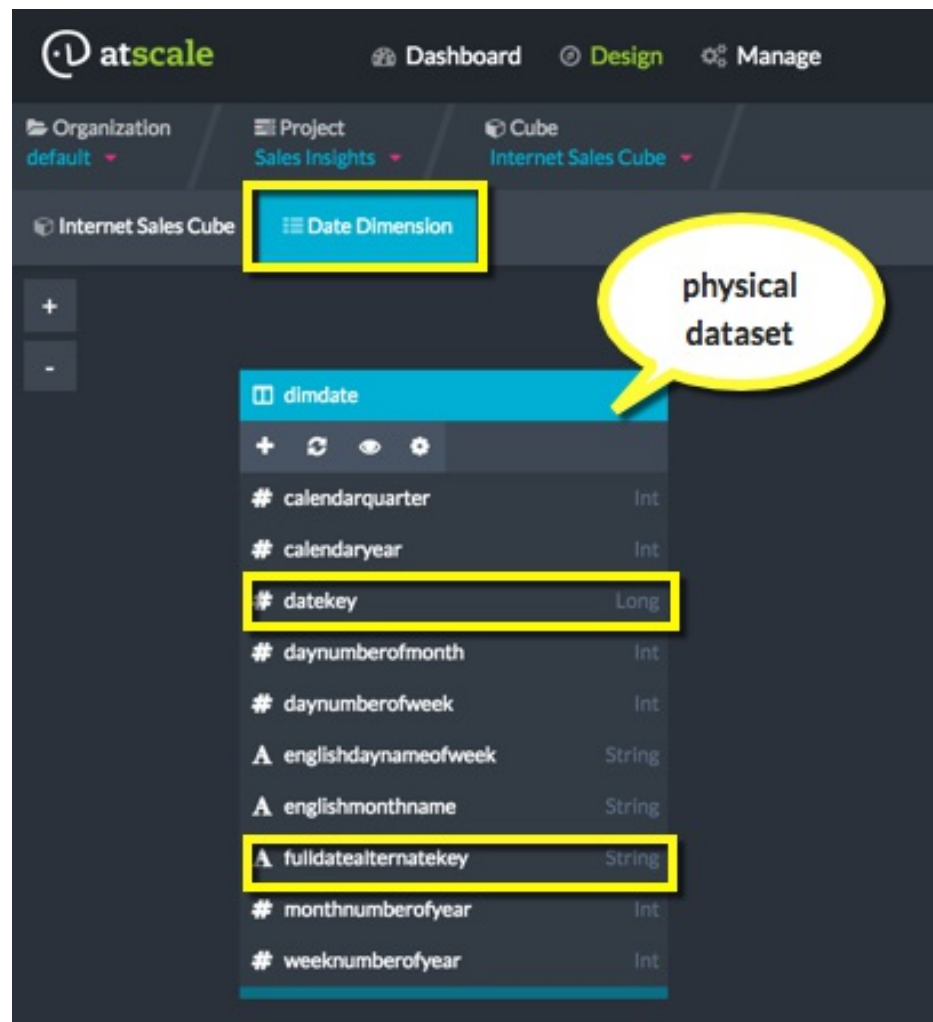



About Keys And Foreign Keys

Two datasets can be joined together on columns that they have in common. The column or columns in the dimension dataset are modeled as the key level of a dimension hierarchy. This key level can then be linked to the corresponding columns in the fact dataset (known as the foreign key).

For example, in this date dimension dataset - there are two possible key columns - one is the date in Unix timestamp format (`LONG`) and the other is the date in the MM-DD-YYYY format (`STRING`).



When you model a logical dimension in AtScale, the dataset's key is always associated to a level in a hierarchy. The key level of a hierarchy is the most granular level. Dimensions with multiple hierarchies can share the same key level if needed. For example, this logical Date dimension has two hierarchies, but the key level (Day) is duplicated and shared between them.



When you create a relationship between a dimension and a fact dataset, the fact dataset must have a corresponding foreign key column (or columns in the case of a compound key) that match to the dimension's key.

These key and foreign key columns must be of the same data type and format for the join to succeed. For example, if one has dates in the format of 01/31/2015 (a string) and one has dates in the format of 1422691200000 (a number), then the join would fail because the data types and formats do not match. You may need to create calculated columns to cast or reformat the data values before defining relationships on them. See [Viewing a Sample of Dataset Rows](#) before creating relationships to make sure the column data types and values sync up between the datasets you plan to join.

