


Many-To-Many Relationships: Multi-Valued Dimensions

A multi-valued dimension is when a fact dataset row can refer to more than one row in a dimension dataset. In AtScale, this is modeled by defining a dimensional bridge or junction table to resolve the many-to-many relationship.

For example, a customer may have many reasons for making a purchase - while the customer-to-purchase relationship is one-to-many, the customer-to-order reason relationship is many-to-many.

To model this in AtScale, a bridge dimension is created between the fact (purchases) and the multi-valued dimension (order reason). The bridge table is considered the multi-valued dimension in this case. This is because the level used to join to the secondary dimension (order reason) is not unique. The relationship between the bridge and order reason is many-to-many.

The bridge table makes it possible to analyze attributes of the customer dimension along with order reason. This is because the bridge table resolves the difference in granularity of the fact data. Customer information is at the granularity of an order (one customer makes an order). However, order reason is at the granularity of an order line item (a customer can have multiple line items in an order, and therefore multiple reasons for purchasing the items in an order).



About Multi-Valued Dimension Attributes

When talking about dimensional multiplicity, a dimension dataset may also have multi-valued attributes. This example of a multi-valued dimension occurs when a dimension row needs to capture multiple values for a single attribute. For example, a customer dataset might have an `interests` column containing a comma-separated list of customer interests (cooking, travel, sports, etc.). The number of values in the array may be arbitrary - a customer could have 0,1,2,3 or 50 interests - presenting a challenge for modeling this type of dimension in a relational model.

Customer ID	Name	Gender	Birthdate	Zip+1e	Interests
C1	Joe Smith	M	3/26/79	90210	Golf, Travel, Cooking
C2	Sarah Jones	F	9/24/80	87631	Travel, Sailing, Art, Crafts
C3	Lisa Nelson	F	4/30/86	10127	Travel
C4	Bill Thompson	M	10/29/59	90210	Photography

Another example of a multi-valued attribute is a single column that contains a map of key:value pairs. AtScale currently supports extracting key:value pairs from map type columns to define new computed columns in a dataset. See [Extract Values from a Map](#).