

Design Tip #132 Kimball Forum Update

By Warren Thornthwaite

It's been two years since we launched the [Kimball Forum](#). At the time, we intended it to be community-driven for practitioners interested in dimensional modeling and other Kimball techniques. It has since grown to nearly 2,000 registered users with a steady level of activity, averaging 43 new topics and 205 posts each month. The following table shows the most viewed topics.

<u>Rank</u>	<u>Replies</u>	<u>Topic</u>
1	38	Joining two Fact tables???
2	23	Bridge Table and Customer Hierarchy
3	23	Regarding Factless Fact Tables
4	22	Reasons to use a 3NF design over a Dimensional Model design for an EDW - Discussion
5	21	Volume and Weight in Same Fact Table
6	20	Enormous data size
7	19	Multiple Fact table with different grain w/conformed dimensions - BO Challenge
8	19	Allocating -Parent child relationship
9	19	Bill Cycle Modeling
10	17	Primary Key of Fact Table

I've been reading through the posts over the last few days, and I'm totally impressed with the quality, detail, and generally helpful nature of the content. The Forum is providing exactly the kind of support we had hoped it would. Let me summarize a few postings from the Forum to give you a sense for what takes place.

[Volume and Weight in Same Fact Table](#)

The person who posted this topic is working on a model to track recycled materials such as oil and tires. These materials are measured in different ways, such as quantity, pounds, and gallons. Measures for some materials include more than one unit of measure. Oil, for example is measured in both gallons and pounds; tires are measured in both quantity and pounds. The question is what are some of the options for modeling multiple measures? The replies offer several options including using a column for each measure, or, using a single measure column with a UnitOfMeasure dimension. Other posters explore the pros and cons of these choices in terms of the size of the fact table, and impact on ease of use and understandability.

[How to load a Slowly Changing Dimension Type 2 with one SQL Merge statement in Oracle](#)

This post is in response to Design Tip #107 which describes how to use the SQL Server MERGE statement to handle slowly changing dimensions. The poster starts with a link to an article describing how to implement an Oracle version of the SQL Server MERGE-based SCD code. A subsequent post offers a 200+ line stored procedure that is a slowly changing dimension processing SQL code generator based on user input such as the source table and key column, target table and key column, and a list of type 1 and type 2 columns. I haven't tried it, but other posters had positive feedback. Of course, your ETL tool may already provide SCD processing capability, but I appreciate the creativity and willingness to

share!

[MDX \(or something else\) for getting history based on last dimension characteristic for SCD](#)

The poster of this topic was originally looking for query guidance on how to total up historical facts based on the current value of a type 2 attribute called Policy Type. Further exploration revealed a need to sum history up to the attribute value at any point in time, not just to the current attribute value. There are several good suggestions, and the poster is now off doing some testing.

If you have a dimensional modeling, ETL, or other DW/BI system problem that you could use some guidance on, the Kimball Forum is a great place to go. It's also a good learning tool, and a way for you to share your own experience and wisdom with the broader Kimball community. Keep posting!