

Design Tip #109 Dos and Don'ts on the Kimball Forum

By Joy Mundy

You've asked (and asked, and asked), and we're finally doing it: we're launching the Kimball Forum! It has gone live already; you can access it from a link right on the [Kimball Group's home page](#). We hope the Kimball Forum will foster a community of dimensional data warehouse practitioners. Many of you have years of experience in business intelligence, and we hope you'll find this a good opportunity to help your colleagues. We at Kimball Group will monitor the forum, we'll chime in sometimes, but we're hoping to develop a vibrant community rather than one that we dominate.

When you first connect to the forum, please take a moment to review the rules. They are summarized here:

- **Keep the discussion polite and civil.** We would rather not delete anyone's posts.
- **Be succinct.**
- **Make an effort to communicate clearly.** If no one can understand your question, no one will answer it.
- **Do not advertise.** That would really annoy us.
- **Don't beg for hand-holding.** This forum is for discussion and analysis, not for answering every question you might dream up. Make an effort to review existing content, both in the Forum and elsewhere on the Kimball website, before posting your question.

Just to get you started on the right foot, here are examples of good questions:

- I am the DW team lead for my company, with a Kimball DW/BI system that's been in production for a few years. We just got a new CIO, who's hired a team of consultants to review all aspects of IT, including the data warehouse. They say our design is inflexible; they say we need an ODS and normalized data warehouse to feed the user-oriented data marts. Can anyone share their experiences to help me argue my position?
- I have a dimension with many alternative rollups or hierarchies. By "many" I mean 20 or so. These rollups are modified, added to, and discontinued fairly frequently (maybe 10 times a year). Although I could make columns in the dimension table for each hierarchy, that design creates problems for ETL. Is there an alternative approach? Although I'm primarily interested in solving this problem for reporting from the RDBMS dimensional data warehouse, I do have some users on OLAP platforms.
- My current business requirements specify that some dimension attributes are Type 1 and some are Type 2. I'm worried about building the system as currently specified, *not* capturing history for Type 1 attributes, and then sometime in the future seeing requirements change. How have others addressed this problem?

Primarily for entertainment purposes, here are some examples of bad questions:

- Attached is the data model for my data warehouse. Can you provide suggestions for improvement?
- What's the difference between SCD-1 and SCD-2?

- I am a master's student, writing my thesis on data warehousing. Can you tell me about it?

Finally, we want to publicly thank Kimball University alum, Brian Jarrett, for initially conceiving and establishing a forum framework that Warren Thornthwaite leveraged for the Kimball Forum. Brian got the ball rolling; Warren kept it in motion.

We hope to see you on the Forum!