

## Kimball Design Tip #16: Hot Swappable Dimensions

By Ralph Kimball

Criterion #18 in the list of Dimensionally Friendly Criteria defines a "hot swappable dimension", which is a dimension with two or more alternative versions. If the dimension is hot swappable, any of the alternative versions of the dimension can be chosen at query time.

There are a number of situations where alternative versions of the same dimension can be very useful. Here are three interesting situations:

- 1) An investment banking house makes available to its clients a large fact table that tracks stocks and bonds on a daily basis over a several year period. The "investment dimension" in this fact table provides information about each stock and bond. But this investment dimension is customized to each client accessing the fact table so that they can each describe and group the investments in interesting and proprietary ways. The different versions of the investment dimension may be completely different including incompatible attribute names and different hierarchy schemes. All clients use the same fact table (hence it only needs to be stored in one place) but each client uses their own investment dimension table as the basis for analyzing the price movements of the stocks and bonds. Viewed from the database server, the clients are busily hot-swapping the investment dimension with each query.
- 2) A retail bank creates a single large fact table that records the month end balances of all the account types in the bank, including checking, savings, mortgage, credit card, personal loans, small business loans, certificates of deposit, student loans, and others. This is a classic case of "heterogeneous products" because the detailed descriptions of each of these account types are wildly different. There is no single description template that can adequately deal with the complexities of all these account types. Therefore we build a simplified account dimension that is meant to join to all the accounts uniformly. We use this simplified account dimension when we are doing cross-selling and up-selling analyses and are looking at the overall portfolio of a customer. But when we restrict our attention to a single account type (e.g., mortgages), we swap in a drastically wider (more fields) dimension that only contains mortgage related attributes. We can do this when we are confident that we have restricted the analysis to just one kind of account. If we have 20 lines of business, we have 21 account dimensions: 1 simplified dimension describing all the accounts and 20 extended dimensions describing disjoint sets of similar accounts.
- 3) A manufacturer wishes to make its shipments fact table available to its trading partners, but needs to shield the orders of the partners from each other. In this case each partner gets their own version of the partner dimension with only their own name appearing in plain text. All the other partners show as "OTHER". Additionally, a mandatory weighting factor field in the dimension is set to 1 for the intended partner and is set to 0 for all others. This weighting factor is uniformly multiplied against all facts in the fact table. In this way, a single shipments fact table can be used to support competitive trading partners in a secure way.

Hot swapping dimensions is straightforward in a standard relational database since the joins between tables can be specified at query time. But if referential integrity is required between the

dimension tables and the fact table, then every swappable version of the dimension must contain the full key set and hence the full set of dimension records. In this case if the swappable dimension is being used to restrict the access to the fact table (as in examples 2 and 3), the restricted rows of the dimension table must contain dummy or null values.

Hot swapping of dimensions is more of a challenge for OLAP systems where the identity of the dimension is built deeply into the fabric of the OLAP data cube. To see how Cognos (PowerPlay) and Microsoft (Analysis Services) handle this hot swappable dimension criterion in their OLAP products, see their white papers at [www.ralphkimball.com/html/dimension.html](http://www.ralphkimball.com/html/dimension.html).