

# Flexible Schema

One of important requirements is to allow schema flexibility: it should be possible to add new columns to existing tables without paying huge price (like rebuilding entire Source table).

We are currently considering two approaches:

- pre-create some number of extra columns in the biggest tables, and rename the columns as they are needed. Periodically a full 'ALTER TABLE' could be run in an organized fashion to add more columns as needed. For example, for Object table (300 columns): 20 extra DOUBLE columns and 10 extra FLOAT would increase Object table by <5% (we wouldn't maintain indexes on the empty columns until they are used).
- after each release, in association with each big table in the unreleased catalog, maintain a narrow table where new columns will be added. This would imply a join between the core big table and the extra table. The extra table would be "reset" each time we cut a new release.

If we allow any user to add columns, it is a security issue. There are also potential access control issues with allowing users to see each other's columns. To resolve these, we may need to maintain one extra narrow table per user. In that case user would need to go through 2 joins.