

The past few weeks have been quite busy at the company and trying to fit the DB tuning into the workflow was quite challenging. I am still going to continue the real-world DB tuning project but compromises had to be made.

The biggest changes are:

- 1) I will be working on two databases instead of just MAIN. In an effort to delineate our production system from the rest of our system even further we are slowly moving tables that were originally in MAIN onto the production system's own DB. This means some of the tuning enhancements will be done on tables that have migrated over to the production environment.
- 2) I will be forgoing database restructuring on a large scale. Since restructuring schemas and tables impacts all systems using them, I will only modify the tables that I am directly tuning. I.e. the overall structure of the DB will remain the same but some tables may change (such as splitting into two, merging into one, creating views, etc.)

Overall, the project remains mostly the same as the initial proposal on the tuning side.

The tuning I proposed such as indexing, materialized views, normal form, and SQL query rewriting will still be performed. I already have one successful enhancement. I discovered a query that took 2-3 seconds and was performed a large number of times in a short period (tens of thousands of similar queries over the span of an hour or two). Realizing that the issue stemmed from the lookup of very specific rows each time, I played around with indexing until I found one that worked (indexing on two columns with one being indexed in descending order). Out of all the indices I tried, this one was the only one to work, lowering the query time to half of a second.