

DataBaseDevContrib

Evidence

For db upgrades we have a unique file, named tiki_db1.xto1.y.sql, with all changes since last release. Today Zaufi said on IRC, "hey what a mess what can we do it's horrible to have to edit that file each time or to make diffs" (translated from russian). Albrown, paradime and I agreed heartfully.

Proposal

Here is a simple solution that can be applied. I submit it to the approbation or contestation of other developpers, so it could become a common and strong rule.

For each modification on the database any of you wish to make, you already change the inter-release file. Now, it could be great if you also commit another file, a patch that should be applied in mysql directly, with only the differences from the trunk. That would make some files, one per CVS version, but would be easier to track and to apply. Even if there are 20 files to apply it can be automatic with one line of shell.

Name convention for the patch could be tiki_patch_1.xx.sql with 1.xx being the version of the current revision of the main db file in the CVS. For example, you patch the file tiki_1.6to1.7.sql with a new table, some inserts, and some table changes. You commit, and you note that new version is 1.50. Then you make a patch with only your new table, some inserts, some alter tables, and data translation, if required. The patch should be named tiki_patch_1.50.sql. The goal is to reach a revision number equivalent to the trunk.

Here is what can be a convention in our development group.
You comments are most welcome.

mose,
with albrown, zaufi and paradime

Mose,

Great idea. This is how I do it at work.

-Ross