
Subject: Re: PostgreSQL Support Classes [Experimental]

Posted by [mirek](#) on Mon, 16 Apr 2007 17:44:42 GMT

[View Forum Message](#) <> [Reply to Message](#)

unodgs wrote on Mon, 16 April 2007 06:13Quote:Obviously, not as sexy as simply puttin more "PRIMARY_KEY" columns, OTOH at least Oracle SQL syntax is similar - you have to add it in single constraint, not as several "PRIMARY_KEY" comlums.

If databse is SQL 92 compatible there is possibility to define primary key in single line. It seems that most of db engines support it.

```
create table t
(
  id integer,
  ver integer,
  type integer,
  ...,
  primary key(id, ver, type)
)
```

Sure, above is a direct equivalent of what have now. SQL allows to define single-column keys in the column and multi-column keys as additional constraint...

Quote:

I proposed putting PRIMARY_KEY attribute next to every column to avoid several macros. Of course I don't know if it is possible to implement it using c preprocessor.

You would have radically change schema code, but I guess it is possible somehow.

Quote:

If not - I suggest to get rid of primary_key attribute as well as DUAL_PRIMARY_KEY and define new macro

PRIMARY_KEY(...) which allow you to define primary key for maximum 5 columns.

Well, but C++ does not have variadic macros (yet?).

BTW, I have just checked and I have used DUAL_PRIMARY_KEY once in my "grand IDIS DB schema" of 500 tables... (OTOH, it is also true that I perhaps deliberately avoid it).

My suggestion is to use current way for now; maybe we can do better when we have all RDBMs implemented and have some experience with it; maybe we will then develop some "common .sch language" to have single .sch for all of them.

Mirek
