

QGIS Application - Bug report #1535
qgspostgresprovider choosing non-unique column as primary key

2009-02-09 03:15 PM - jcs -

Status:	Closed	
Priority:	Low	
Assignee:	Jürgen Fischer	
Category:	Data Provider	
Affected QGIS version:		Regression?: No
Operating System:	All	Easy fix?: No
Pull Request or Patch supplied:		Resolution: fixed
Crashes QGIS or corrupts data:		Copied to github as #: 11595
Description		
Table A (rid primary key, bid int, loc geometry)		
Table B (rid primary key, id2 unique)		
View V is (select Table.rid, Table B.id2, Table A.loc from Table A left outer join Table B on Table A.bid=Table B.rid;)		
The problem is Qgis may select Table B.id2 as its primary key because of the unique constraint in the table definition even though V.id2 is not constrained to be unique in the view.		
For example if Table A has		
rid bid loc		
1 1 POINTA		
2 2 POINTB		
3 1 POINTC		
and Table B has		
rid id2		
1 47		
2 48		
View V would have		
rid id2 loc		
1 47 POINTA		
2 48 POINTB		
3 47 POINTC		
Resolution... columns obtained through a left outer join should not be included for consideration as a primary key... this could be implemented by a modification to findColumns() in qgspostgresprovider to not return such columns.		

History

#1 - 2009-02-09 04:45 PM - jcs -

after a little more code diving...

the above case will be caught by `uniqueData()` if and only if the tables are initialized before the layer is created. if in a dynamic situation where the tables are initialized after the layer is created, the error condition still exists.

i still think a better resolution is to catch the join and rule out any of those columns as you cannot be sure they will not one day become non-unique.

#2 - 2009-02-10 08:51 AM - jcs -

My attached bugfix allows the user to specify which column they would like to use as the primary key by filling out the added "key" data member of the `[[QgsDataSourceURI]]`.

The `[[QgsPostgresProvider]]` will do some verification on the specified column if it exists, otherwise it will do the current method of trying to find a column to use.

#3 - 2009-02-19 10:50 PM - gjm -

I think that we also need to provide this sort of UI functionality:

- when selecting tables to load, allow the user to indicate that he/she wants to select the id column manually (and be prompted for it)
- when no suitable columns are found automatically, prompt the user to choose a suitable column

In both cases, qgis should continue to check that the column contains unique data.

#4 - 2009-04-26 03:53 PM - Jürgen Fischer

in commit:468550c6 (SVN r10658) the key column of postgres view layers is now stored to the project file and retried first on reload of the project.

#5 - 2009-04-27 07:34 AM - jcs -

Replying to [comment:6 jef]:

| *in commit:468550c6 (SVN r10658) the key column of postgres view layers is now stored to the project file and retried first on reload of the project.*

in my situation, i am using the qgis library and neither the application directly nor a project file. also, the problem i was having was preventing me from loading the layer initially. it would seem, then, that it would never make it into the project file?

#6 - 2009-04-27 09:23 AM - Jürgen Fischer

Replying to [comment:7 jcs]:

| *Replying to [comment:6 jef]:*

|

in commit:468550c6 (SVN r10658) the key column of postgres view layers is now stored to the project file and retried first on reload of the project.

in my situation, i am using the qgis library and neither the application directly nor a project file. also, the problem i was having was preventing me from loading the layer initially. it would seem, then, that it would never make it into the project file?

Right. The original problem is probably still there. AFAICS your patch doesn't address that issue either. So I didn't close the bug.

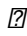
```
<pre>
setDataSourceUri( mUri.uri() );
</pre>
```

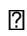
after locating the view's key column, when there is no give column or if it turned out to be unsuitable. Which causes the key to actually make it into the project file.

That avoids the need to go through the whole expensive key column lookup procedure again on reload and should solve the problem Andreas was reporting in <http://lists.osgeo.org/pipermail/qgis-user/2009-April/005182.html>

#7 - 2009-04-27 10:37 AM - jcs -

Replying to [comment:8 jef]:


Right. The original problem is probably still there. AFAICS your patch doesn't address that issue either. So I didn't close the bug.


That's correct, I gave up on the bug and took the easy way out :). Something tells me that no matter what algorithm is used to look for a unique column, a view could be constructed that breaks it. That's just my pessimistic hunch, I have no proof.

```
>
<pre>
> setDataSourceUri( mUri.uri() );
</pre>
>
> after locating the view's key column, when there is no give column or if it turned out to be unsuitable. Which causes the key to actually make it into the project file.
```

This rev looks great, thanks for looking into this.

#8 - 2009-07-13 10:54 AM - Paolo Cavallini

Can this be considered a solution to the problem? Should the ticket be left open, or can be closed?

#9 - 2009-07-13 11:45 AM - jcs -

Replying to [comment:10 pcav]:



Can this be considered a solution to the problem? Should the ticket be left open, or can be closed?



I think it can be closed.

#10 - 2009-07-13 12:12 PM - Jürgen Fischer

- Resolution set to fixed

- Status changed from Open to Closed

Replying to [comment:11 jcs]:



Replying to [comment:10 pcav]:



Can this be considered a solution to the problem? Should the ticket be left open, or can be closed?



I think it can be closed.



Ok then.

Files

bug1535fix.diff	6.16 KB	2009-02-10	jcs -
-----------------	---------	------------	-------