

QGIS Application - Bug report #20109

Editing Oracle Spatial view in QGIS

2018-10-15 11:30 AM - Maxime Resibois

Status:	Closed	
Priority:	Normal	
Assignee:	Jürgen Fischer	
Category:	Data Provider/Oracle	
Affected QGIS version:	3.3(master)	Regression?: No
Operating System:	Windows 7	Easy fix?: No
Pull Request or Patch supplied:	No	Resolution:
Crashes QGIS or corrupts data:	No	Copied to github as #: 27931
Description		
<p>I try with QGIS to insert new features in a multi-table Oracle Spatial view that is updatable by adding an "instead of" trigger. The view is "key-preserved" and I don't have any problem when I insert a new feature directly in SQL.</p> <p>But when I try to do this in QGIS (QGIS 2.18.24 as QGIS 3.2) I get this error :</p> <p>Oracle error while adding features: Oracle error: Could not retrieve feature id -2 SQL: Error: SELECT "OBJECTID" FROM "MY_SCHEMA"."MY_VIEW" WHERE ROWID=:f</p> <p>Despite this error, the feature is inserted in my database anyway. The problem is, if I save layer edits another time, my feature would be recorded 2 times.</p> <p>I think the problem is there are no ROWID pseudocolumns in oracle views while there are in tables and QGIS try to use this pseudocolumn.</p>		

Associated revisions

Revision 855b3b4e - 2019-01-22 03:47 PM - Jürgen Fischer

oracle provider: check for valid lastInsertId() on ::addFeatures (fixes #20109)

Revision 3d4c6c4e - 2019-01-23 01:58 PM - Jürgen Fischer

oracle provider: check for valid lastInsertId() on ::addFeatures (fixes #20109)

(cherry picked from commit 855b3b4e26377647de5c5f9d38485f9f00bc9257)

History

#1 - 2018-10-15 10:15 PM - Giovanni Manghi

- Status changed from Open to Feedback
- Priority changed from High to Normal

Does it works the same way in QGIS 3.2/3.3?

#2 - 2018-10-16 11:45 AM - Maxime Resibois

Yes it works the same way in QGIS 3.2 and 3.3 nightly

#3 - 2018-10-16 02:15 PM - Giovanni Manghi

- Status changed from Feedback to Open
- Affected QGIS version changed from 2.18.24 to 3.3(master)

#4 - 2018-10-16 05:23 PM - Maxime Resibois

A small precision : this issue occurs only with multi-table view. With a view that is a mirror of an unique table there is no problem.

#5 - 2018-10-17 10:47 AM - Maxime Resibois

Ok I've continued my tests. The issue always occurs when I have an "instead of" trigger on my view even if it's a one table view.

#6 - 2018-11-09 08:19 PM - Ivan Lucena

- File *qgsoracleprovider.cpp.patch* added

I was looking at the code on *qgsoracleprovider.cpp* I decided to makes some changes but I would like to confirm if they are safe.

The variable "QSqlQuery getfid" is used to prepare a query based on ROWID, even though it is inside a block where *mPrimaryKeyType* is not *PktRowId*:

[a] [[<https://github.com/qgis/QGIS/blob/30cf2d37bc6b66a75e4313dc2752cb5e0de3e3c8/src/providers/oracle/qgsoracleprovider.cpp#L1243>]]

After the insert statement is executed there is a sequence of code that will run the query prepared on *getfid*:

[b] [[<https://github.com/qgis/QGIS/blob/30cf2d37bc6b66a75e4313dc2752cb5e0de3e3c8/src/providers/oracle/qgsoracleprovider.cpp#L1331>]]

That code is supposed to update the features ids. But that is also what the following code does (without "getfid"):

[c] [[<https://github.com/qgis/QGIS/blob/30cf2d37bc6b66a75e4313dc2752cb5e0de3e3c8/src/providers/oracle/qgsoracleprovider.cpp#L1365>]]

By eliminating the "getfid" query completely [a,b] this problem seems to be fixed but I am wondering if that is safe.

What was the intention in [a,b]? To query for the auto-increment primary key? If that is the case, can we do the *getfid* query [a] based on primary keys instead of *rowid*?

#7 - 2018-11-09 08:25 PM - Nyal Dawson

Can you open a pull request on GitHub with your patch? It's easier to discuss and review code that way

#8 - 2018-11-11 02:51 PM - Jürgen Fischer

Ivan Lucena wrote:

What was the intention in [a,b]? To query for the auto-increment primary key? If that is the case, can we do the *getfid* query [a] based on primary keys instead of *rowid*?

We need to update the inserted feature with the keys of the new row - and we get it based on lastInsertId() which is the ROWID of the new row.

#9 - 2018-11-12 04:48 PM - Ivan Lucena

Thanks, Jürgen.

If I understand it correctly, on addFeatures(), the feature list that is coming from QGIS editor UI doesn't have an Id to refers to the database table yet.

It is only after the commit is completed successfully that addFeatures() can update the feature list that is the memory.

That would prevent a lot of mistakes, like insert the same feature again, or use the same id as another simultaneous user session.

That works perfectly as long we don't have a view with an "instead of insert" where the logic of the insert is on the custom-made user code on PL/SQL.

And Oracle documentation advise against the use of ROWID on that case.

It seems like one option is to use the primary key as feature Id. That is what they do in QgsPostgresProvider::addFeature(). They even use " RETURNING " to support for auto-increment columns.

Could we do the same for Oracle?

#10 - 2018-11-13 12:51 AM - Jürgen Fischer

<https://github.com/qgis/QGIS/pull/8466>

#11 - 2018-11-14 02:31 PM - Ivan Lucena

Thanks, Jürgen.

I think you are on the right path but I tried your new code with Maxime's data and got this error:

```
2018-11-14T08:12:03 CRITICAL Layer V_TMP_VIEW_EDIT : Oracle error while adding features: Oracle error: Could not insert feature -2
SQL: ORA-22816: unsupported feature with RETURNING clause
Unable to execute statement
Error: INSERT INTO "SCOTT"."V_TMP_VIEW_EDIT"("GEOM","OBJECTID","NAME","DATE_TEST") VALUES (:a,:bb,:bc,:bd) RETURNING
"OBJECTID" INTO :be
2018-11-14T08:12:03 WARNING Commit errors : Could not commit changes to layer V_TMP_VIEW_EDIT
```

I am not sure that we need to return OBJECTID if we are inserting on it. The user provides its value and if it is duplicated or something it will raise a proper error. But maybe there are other use cases that I am not aware of.

#12 - 2018-11-14 05:54 PM - Jürgen Fischer

Ivan Lucena wrote:

| *I think you are on the right path but I tried your new code with Maxime's data and got this error:*

Which data? INSERT...RETURNING works fine for me - but on a plain table - if that's unsupported for VIEW the path is apparently not that right.

2018-11-14T08:12:03 CRITICAL Layer V_TMP_VIEW_EDIT : Oracle error while adding features: Oracle error: Could not insert feature -2
SQL: ORA-22816: unsupported feature with RETURNING clause
Unable to execute statement
Error: INSERT INTO "SCOTT"."V_TMP_VIEW_EDIT"("GEOM","OBJECTID","NAME","DATE_TEST") VALUES (:a,:bb,:bc,:bd) RETURNING
"OBJECTID" INTO :be
2018-11-14T08:12:03 WARNING Commit errors : Could not commit changes to layer V_TMP_VIEW_EDIT

I am not sure that we need to return OBJECTID if we are inserting on it. The user provides its value and if it is duplicated or something it will raise a proper error. But maybe there are other use cases that I am not aware of.

A trigger might update the primary key - eg. when there is a nextval. We cannot even leave the primary key NULL in the GUI anymore as the provider detects the automatic NOT NULL constraint on primary keys.

#13 - 2018-11-17 03:14 PM - Ivan Lucena

- File test_auto_incr.sql added
- File sample.sql added

Replacing ROWID by RETURNING works very well with SEQUENCE/NEXTVAL. See attachment.

I also wrote a script that builds a scenario very similar to the bug use case (sample.sql).

It seems like RETURNING is not going to work on that case. Maybe it would be better to have a master table with geometry columns and build two views, one with and one without geometry.

#14 - 2019-01-22 03:47 PM - Jürgen Fischer

- Status changed from Open to Closed
- % Done changed from 0 to 100

Applied in changeset commit:qgis|855b3b4e26377647de5c5f9d38485f9f00bc9257.

Files

qgsoracleprovider.cpp.patch	1.96 KB	2018-11-09	Ivan Lucena
sample.sql	2.76 KB	2018-11-17	Ivan Lucena
test_auto_incr.sql	1.21 KB	2018-11-17	Ivan Lucena