

QGIS Application - Bug report #20077

geopackage - table operations via network shares take a very long time

2018-10-11 11:31 AM - Burghardt Scholle

Status:	Open	
Priority:	Normal	
Assignee:		
Category:	Data Provider	
Affected QGIS version:	3.3(master)	Regression?: No
Operating System:	Windows 7 64bit, Windows 10	Easy fix?: No
Pull Request or Patch supplied:		Resolution:
Crashes QGIS or corrupts data:		Copied to github as #: 27899
Description I have created a new geopackage with OpenStreetMap data under QGIS Master. Unfortunately, there is also the problem with shared network folders here (samba share). Creating a new table with a name that already existed takes a long time. I create the new table via the context menu "Export -> Save as" of the layer. Once triggered, the process can only be aborted under Windows using the TaskManager. When trying to close QGIS, the message documented in the attached screenshot appears. After a click on "Yes", QGIS is not closed. If I copy the same geopackage into a local directory, the described process can be successfully completed within a few moments. (Note: This bug report results from bug report #20068. Further information can be found there.)		
Related issues: Related to QGIS Application - Bug report # 20068: Geopackage: problems with r... Closed 2018-10-10		

History

#1 - 2018-10-11 11:37 AM - Jürgen Fischer

- Related to Bug report #20068: Geopackage: problems with renaming or drop tables added

#2 - 2018-10-11 12:44 PM - Even Rouault

Can you retry your tests by defining the following environment variables (may require restarting QGIS each time you change) ?

- OGR_SQLITE_JOURNAL=DELETE
- SQLITE_USE_OGR_VFS=YES
- and possibly the combinations of both

Note that SQLite is notoriously known not to be appropriate for databases stored on network shares...

#3 - 2018-10-11 04:41 PM - Burghardt Scholle

Okay, it's been a long day but ...

wow, the environment variable "SQLITE_USE_OGR_VFS=YES" has done it - under Windows and Xubuntu! How did I go about it?

I installed QGIS Nightly (from <https://qgis.org/ubuntu-nightly>) in a fresh Xubuntu installation. This version only comes with GDAL/OGR 2.2.3, but no matter.

For testing purposes I have executed an "export SQLITE_USE_OGR_VFS=YES" from the command line and then started QGIS from there as well.

Loading the QGIS project with my "network"-geopackage and testing the table creation was successful and very fast! Awesome.

After this success I set a global variable by creating a file "qgis-gpkg.sh" in the directory "/etc/profile.d" with the following content:

1. QGIS environment variable for geopackage and shared network folder
export SQLITE_USE_OGR_VFS=YES

The computer must then be restarted. Afterwards, the table operations via network shares worked perfectly and with the expected performance.

Under Windows I added "export SQLITE_USE_OGR_VFS=YES" in my system environment variables. Afterwards I had to log off and on again so that the change would have an effect. Thank you very much for pointing that out!

In view of the fact that geopackage will replace shape as the default format in the future, maybe we should set this system variables during installation in future QGIS versions. I believe that in practice many users use network drives (because of a central data backup, data sharing, etc.).

#4 - 2018-10-11 09:00 PM - Even Rouault

The use of SQLITE_USE_OGR_VFS=YES has the side effect of not implementing SQLite locking, which explains the speed increase. But this may lead to database corruption in case of concurrent edits. So this is not the perfect solution. But that might be something to enable when the database is detected to be on a network share

#5 - 2018-10-12 07:16 AM - Burghardt Scholle

Okay, that's not nice, of course. The concurrent use is one of the main advantage when storing data on a shared network folder. What surprises me is that the problems only occur when one assign a table name that already existed in the geopackage. If one create a new table with an unused name, these problems do not occur. With QGIS master, a table with a previously used name is created correctly - it only takes forever. I don't understand why there is such a huge difference between a geopackage stored locally or in the network.

Files

GeoPackage_1.png	11.8 KB	2018-10-11	Burghardt Scholle
------------------	---------	------------	-------------------