

QGIS Application - Bug report #19968

Very long loading times when opening the same FileGDB the second time

2018-09-27 03:12 PM - Casper Børgesen

Status:	Open	
Priority:	Normal	
Assignee:		
Category:	Data Provider/OGR	
Affected QGIS version:	3.3(master)	Regression?: No
Operating System:	Windows 7/10	Easy fix?: No
Pull Request or Patch supplied:		Resolution:
Crashes QGIS or corrupts data:		Copied to github as #: 27790
Description		
<p>The problem and how to reproduce:</p> <ol style="list-style-type: none">1. I have a FileGDB which I drag into QGIS (14 seconds).2. I select all layers and press OK (7 seconds).3. I remove all the layers (empty layer list and this process is slow too).4. I drag the same FileGDB into QGIS again (14 seconds).5. I select all layers and press OK (I was disturbed by a colleague after about 2.5 minutes and all layers was loaded when I came back after 15 minutes). <p>As seen on the above timings loading the same FileGDB the second time is magnitudes slower than the first time around.</p> <p>All I can see is that a lot of locking/unlocking of the files GDB_Items and GDB_ItemTypes during all the processes.</p> <p>It might be related to issue #18342 where Even Rouault did some optimization.</p> <p>I have attached the test data set used for the timings above (the same as used in issue #18342).</p> <p>My tests was done on a Laptop with an SSD hard-drive but the issue is also present on my workstation.</p> <p>I have installed QGIS nightly using OSGeo4W x64 and the package gdal-filegdb is also installed.</p> <p>I have ArcMap 10.6 and ArcGIS Pro 2.2 installed.</p>		

History

#1 - 2018-09-27 03:21 PM - Casper Børgesen

Casper Børgesen wrote:

1. I have a FileGDB which I drag into QGIS (14 seconds).
2. I select all layers and press OK (7 seconds).
3. I remove all the layers (empty layer list and this process is slow too).
4. I drag the same FileGDB into QGIS again (14 seconds).
5. I select all layers and press OK (I was disturbed by a colleague after about 2.5 minutes and all layers was loaded when I came back after 15 minutes).

I wanted to see if ESRI had the same problem and I repeated the test in ArcGIS Pro:

- Step 1 and 2 is done in about 14 seconds.
- Step 3 is done almost instantly.
- Step 4 and 5 is also done in about 14 seconds.

QGIS seems to keep a lock on all feature types when the FileGDB is open, where ArcGIS Pro doesn't. I think that could be the cause to why QGIS is slower when removing all the layers.

#2 - 2018-09-27 05:17 PM - Giovanni Manghi

- Status changed from Open to Feedback

Does it work differently in QGIS 2.18?

#3 - 2018-09-28 11:24 AM - Casper Børjesen

- Status changed from Feedback to Open

Giovanni Manghi wrote:

| Does it work differently in QGIS 2.18?

I have just tried on the two machines and QGIS 2.18 has the issue already in step 2.

Would it help if I recorded a video of what I am doing along with all the file activity in the GDB folder?

#4 - 2018-09-28 05:08 PM - Giovanni Manghi

- Priority changed from High to Normal

- Regression? changed from Yes to No

Casper Børjesen wrote:

| Giovanni Manghi wrote:

| Does it work differently in QGIS 2.18?

| I have just tried on the two machines and QGIS 2.18 has the issue already in step 2.

so not a regression.

| Would it help if I recorded a video of what I am doing along with all the file activity in the GDB folder?

yes please.

#5 - 2018-10-01 08:32 AM - Casper Børjesen

Giovanni Manghi wrote:

| yes please.

I have uploaded a short video to youtube where you can see whats happening in QGIS and in the FileGDB folder when loading and removing the FileGDB from QGIS.

<https://www.youtube.com/watch?v=hvcPVObfBFA>

#6 - 2018-10-02 10:30 AM - Giovanni Manghi

- Status changed from Open to Feedback

Just tested on master/linux and the operations as you show them in the screencast are basically instantaneous.

#7 - 2018-10-02 12:17 PM - Casper Børjesen

- Status changed from Feedback to Open

Giovanni Manghi wrote:

| Just tested on master/linux and the operations as you show them in the screencast are basically instantaneous.

I am only on Windows so I can't provide feedback regarding Linux.

#8 - 2018-10-03 08:32 PM - Casper Børjesen

Giovanni Manghi wrote:

| Just tested on master/linux and the operations as you show them in the screencast are basically instantaneous.

I just uninstalled the OGR FileGDB driver and reran the tests. Without the OGR FileGDB driver I probably get the same instantaneous loading like you do.

I guess that you don't have the OGR FileGDB driver installed (since I expect it to be Windows only?).

I have tried the same tests with and without the OGR FileGDB driver installed on my private PC (which isn't infected by any ESRI products) and I get the same results, with the driver its very slow and without its very fast. In most cases I need the write capability of the OGR FileGDB driver and I haven't found an easy way to select which driver to use - besides (un)installing :s

So I think this issue is related to the OGR FileGDB driver (or how it is used by QGIS).

#9 - 2018-10-04 10:35 AM - Giovanni Manghi

Casper Børjesen wrote:

Giovanni Manghi wrote:

| Just tested on master/linux and the operations as you show them in the screencast are basically instantaneous.

I just uninstalled the OGR FileGDB driver and reran the tests. Without the OGR FileGDB driver I probably get the same instantaneous loading like you do.

I guess that you don't have the OGR FileGDB driver installed (since I expect it to be Windows only?).

I have tried the same tests with and without the OGR FileGDB driver installed on my private PC (which isn't infected by any ESRI products) and I get the same results, with the driver its very slow and without its very fast. In most cases I need the write capability of the OGR FileGDB driver and I haven't found an easy way to select which driver to use - besides (un)installing :s

So I think this issue is related to the OGR FileGDB driver (or how it is used by QGIS).

I just tested with both the proprietary "ogr filegdb" (esri fliegdb) available to be installed manually with osgeo4w and the open source one (openFileGDB) is installed by default (as is part of ogr) on any operating system when installing QGIS. On Windows 10/QGIS 3.2 (osgeo4w) I can't notice any delay or long wait when adding/removing/re-adding your datasources.

Have you tried on a clean profile?

#10 - 2018-10-04 01:39 PM - Casper Børjesen

Giovanni Manghi wrote:

Have you tried on a clean profile?

What do you mean by a *clean profile*?

#11 - 2018-10-04 01:41 PM - Giovanni Manghi

Casper Børjesen wrote:

Giovanni Manghi wrote:

Have you tried on a clean profile?

What do you mean by a clean profile?

In QGIS3 you can have multiple profiles (settings > user profiles). Testing in a new profile is necessary to exclude issues caused by 3rd party plugins and/or legacy configurations.

#12 - 2018-10-05 02:23 PM - Casper Børjesen

Giovanni Manghi wrote:

Casper Børjesen wrote:

Giovanni Manghi wrote:

Have you tried on a clean profile?

What do you mean by a clean profile?

In QGIS3 you can have multiple profiles (settings > user profiles). Testing in a new profile is necessary to exclude issues caused by 3rd party plugins and/or legacy configurations.

I have created a new profile and using that profile I still get the same slow results :(

#13 - 2018-10-05 03:52 PM - Andrea Giudiceandrea

I can reproduce the issue on Windows 7 64 bit (core i5, 8GB) with QGIS 3.3.0-Master (3032af2dcc) (qgis-dev-3.3.0-84) 64 bit.

The issue only happens when the testdata_2.gdb directory (extracted from the provided zip file) is dragged and dropped in the map or when the layers are added through "Add vector layer | Data Source Manager | Vector" / "Source Type: Directory" and only if the GDAL/OGR FileGDB driver plugin (gdal-filegdb 2.2.4-2) is installed (regardless of the Source format chosen among "OpenFileGDB" and "ESRI FileGDB").

I can also confirm that the slowdowns seem related to the creation of lock files in the testdata_2.gdb directory during the import/removal operations.

It doesn't happen if the GDAL/OGR FileGDB driver plugin is not installed.

It doesn't happen if the provided testdata_2.gdb.zip file (not the extracted testdata_2.gdb directory) is dragged and dropped into the map (regardless the GDAL/OGR FileGDB driver plugin is installed or not).

#14 - 2018-10-06 10:52 AM - Giovanni Manghi

- Status changed from Open to Feedback

Andrea Giudiceandrea wrote:

I can reproduce the issue on Windows 7 64 bit (core i5, 8GB) with QGIS 3.3.0-Master (3032af2dcc) (qgis-dev-3.3.0-84) 64 bit.

The issue only happens when the testdata_2.gdb directory (extracted from the provided zip file) is dragged and dropped in the map or when the layers are added through "Add vector layer | Data Source Manager | Vector" / "Source Type: Directory" and only if the GDAL/OGR FileGDB driver plugin (gdal-filegdb 2.2.4-2) is installed (regardless of the Source format chosen among "OpenFileGDB" and "ESRI FileGDB").

weird, I'm pretty sure I tested this scenario on a clean environment (VM snapshot) and I could not confirm. Anyway, if this observation is right then the title and description of this ticket should be changed to better reflect the source of the issue.

#15 - 2018-10-08 10:02 AM - Casper Børgesen

- Status changed from Feedback to Open

Giovanni Manghi wrote:

... Anyway, if this observation is right then the title and description of this ticket should be changed to better reflect the source of the issue.

I am not allowed to change anything related to the question or subject.

#16 - 2018-10-08 11:20 AM - Giovanni Manghi

Casper Børgesen wrote:

Giovanni Manghi wrote:

... Anyway, if this observation is right then the title and description of this ticket should be changed to better reflect the source of the issue.

I am not allowed to change anything related to the question or subject.

suggest a new title and I will change it afterwards.

#17 - 2018-10-08 01:01 PM - Casper Børgesen

Giovanni Manghi wrote:

Casper Børgesen wrote:

Giovanni Manghi wrote:

... Anyway, if this observation is right then the title and description of this ticket should be changed to better reflect the source of the issue.

I am not allowed to change anything related to the question or subject.

suggest a new title and I will change it afterwards.

Title: Slow performance when using OGR FileGDB driver to read file geodatabases.

Description:

When loading file geodatabases into QGIS there is a big performance difference between using the OpenFileGDB driver and the OGR FileGDB driver. Opening geodatabases using the OpenFileGDB driver is almost instantaneous where using the OGR FileGDB driver seems to spend a lot of time opening and closing .lock files during the import/removal operations.

- It affects QGIS 2.18, 3.2 and MASTER.
- It seems to be unaffected by whether any ESRI tools are installed or not.
- It might not affect *Linux* systems.

It might be related to issue #18342 where Even Rouault did some optimizations.

Steps to reproduce (<https://www.youtube.com/watch?v=hvcPVObfBFA>):

The OGR FileGDB driver needs to be installed in order to reproduce and the *testdata_2.gdb.zip* can be used in the test.

1. I have a FileGDB which I drag into QGIS (14 seconds).
2. I select all layers and press OK (7 seconds).
3. I remove all the layers (empty layer list and this process is slow too).
4. I drag the same FileGDB into QGIS again (14 seconds).
5. I select all layers and press OK (I was disturbed by a colleague after about 2.5 minutes and all layers was loaded when I came back after 15 minutes).

#18 - 2018-10-09 10:44 AM - Giovanni Manghi

Title: Slow performance when using OGR FileGDB driver to read file geodatabases.

should not reflect the finding from here?
it happens, but only in some circumstance, correct?

#19 - 2018-10-09 11:05 AM - Casper Børgesen

Giovanni Manghi wrote:

Title: Slow performance when using OGR FileGDB driver to read file geodatabases.

should not reflect the finding from here?

it happens, but only in some circumstance, correct?

If I compare to using the OpenFileGDB driver, then the performance is slower (not as instantaneous as the OpenFileGDB driver) all the time - no matter if I use 2.18, 3.2 or MASTER. The issue seems to be much more severe the second time I load the data set, which is what I try to show in the video.

#20 - 2018-10-09 08:50 PM - Giovanni Manghi

If I compare to using the OpenFileGDB driver, then the performance is slower (not as instantaneous as the OpenFileGDB driver) all the time - no matter if I use 2.18, 3.2 or MASTER. The issue seems to be much more severe the second time I load the data set, which is what I try to show in the video.

so you don't confirm only when the layers are added in the ways as stated here?

#19968-13 (edited)

from that comment I deduce that if the layers are added using the QGIS Browser the issue does not show(?).

#21 - 2018-10-09 11:06 PM - Andrea Giudiceandrea

Giovanni Manghi wrote:

from that comment I deduce that if the layers are added using the QGIS Browser the issue does not show(?).

With the GDAL/OGR FileGDB driver plugin installed, using the Browser panel to add the layers to the map it seems that the second time is not slower then the first time: they are almost equally slow (it's a matter of minutes for add or remove the layers). Without the FileGDB driver plugin, each task takes just a few seconds.

Files

testdata_2.gdb.zip	110 KB	2018-09-27	Casper Børgesen
--------------------	--------	------------	-----------------