

QGIS Application - Bug report #19509

Slow opening of large GPX file on Windows

2018-07-30 11:57 AM - David Fiedler

Status: Closed	
Priority: Normal	
Assignee:	
Category: Data Provider	
Affected QGIS version: 3.2.1	Regression?: No
Operating System: Windows	Easy fix?: No
Pull Request or Patch supplied: No	Resolution: no timely feedback
Crashes QGIS or corrupts data: No	Copied to github as #: 27337

Description

The opening of large GPX file (on my computer, it is noticeable when using files larger than 50MB) is very slow on windows. The problem is not the rendering, the rendering is reasonably fast with respect to file size. The lag is before the rendering process starts, before and after the dialog when we can decide what elements to import from the gpx file. The whole UI freezes for a couple of seconds for 50MB file up to over half hour in case of 3GB gpx file. I've tested the behavior in both 3.2 and 2.18 versions of QGIS.

When opening the same files on Linux, the time before rendering starts are significantly shorter (4 vs 35 mins for 3GB file).

History

#1 - 2018-07-30 12:06 PM - Giovanni Manghi

- Status changed from Open to Feedback

Does it works better on QGIS LTR? Can you point us to the data?

#2 - 2018-07-30 03:54 PM - David Fiedler

If I understand it correctly, version 2.18 is LTR, so yes, I've already tested it on LTR. I didn't benchmarked it, so I don't have exact times, but the freeze was also very long on 2.18. Will it be helpful if I send you exact times for both versions?

Concerning the data, I have to anonymize it first.

#3 - 2018-07-30 03:58 PM - Giovanni Manghi

David Fiedler wrote:

| Will it be helpful if I send you exact times for both versions?

no, not for me at least.

How ogrinfo performs against this large datasource?

#4 - 2018-07-30 06:15 PM - David Fiedler

I've just tested loading the file in ogrinfo and it performs very poorly. If this tool is used in QGIS to perform analysis of the file before the rendering starts, the

slowness is no longer a mystery.

#5 - 2018-07-30 08:46 PM - Giovanni Manghi

David Fiedler wrote:

I've just tested loading the file in ogrinfo and it performs very poorly. If this tool is used in QGIS to perform analysis of the file before the rendering starts, the slowness is no longer a mystery.

it does indeed. This ticket should be closed here and reopened in the gdal tracker.

#6 - 2018-08-10 10:26 AM - David Fiedler

Concerning ogrinfo, which command is relevant to QGIS and therefore I should evaluate it?

#7 - 2018-08-12 01:42 PM - Giovanni Manghi

David Fiedler wrote:

Concerning ogrinfo, which command is relevant to QGIS and therefore I should evaluate it?

QGIS uses OGR to handle a number of formats (while for others it has native providers), so if it is slow in OGR it will also be slow in QGIS. Please post this issue in the GDAL/OGR bug tracker.

#8 - 2018-08-17 11:34 AM - David Fiedler

Giovanni Manghi wrote:

David Fiedler wrote:

Concerning ogrinfo, which command is relevant to QGIS and therefore I should evaluate it?

QGIS uses OGR to handle a number of formats (while for others it has native providers), so if it is slow in OGR it will also be slow in QGIS. Please post this issue in the GDAL/OGR bug tracker.

I understand that, however, I don't know which OGR command is representative, I don't know how ogrinfo is used in QGIS. For example when I type ogrinfo -so FILE, it's quick. Typing ogrinfo -al FILE is slow on both platforms. So I would need a command to benchmark, so far, there is no difference (in speed) between linux and windows version when using ogrinfo (like it is in QGIS).

#9 - 2018-08-17 02:18 PM - Jürgen Fischer

David Fiedler wrote:

I understand that, however, I don't know which OGR command is representative, I don't know how ogrinfo is used in QGIS. For example when I type ogrinfo -so FILE, it's quick. Typing ogrinfo -al FILE is slow on both platforms. So I would need a command to benchmark, so far, there is no

difference (in speed) between linux and windows version when using ogrinfo (like it is in QGIS).

ogrinfo and QGIS's ogrprovider both use OGR functions to fetch the information - but QGIS doesn't directly use ogrinfo, so there is no clear equivalent.

#10 - 2018-08-19 11:04 AM - Giovanni Manghi

ogrinfo and QGIS's ogrprovider both use OGR functions to fetch the information - but QGIS doesn't directly use ogrinfo, so there is no clear equivalent.

should be close this as upstream?

#11 - 2018-08-19 02:33 PM - David Fiedler

Jürgen Fischer wrote:

David Fiedler wrote:

I understand that, however, I don't know which OGR command is representative, I don't know how ogrinfo is used in QGIS. For example when I type ogrinfo -so FILE, it's quick. Typing ogrinfo -al FILE is slow on both platforms. So I would need a command to benchmark, so far, there is no difference (in speed) between linux and windows version when using ogrinfo (like it is in QGIS).

ogrinfo and QGIS's ogrprovider both use OGR functions to fetch the information - but QGIS doesn't directly use ogrinfo, so there is no clear equivalent.

The problem is, that commands I have tried appear to be solved in a similar time using Windows and Linux. In other words, so far I have no evidence that GDAL/ogrprovider is the cause of the slowness. It appeared like the cause first, but then I realized that the command I tried is slow on Linux too. Sorry for the confusion.

#12 - 2018-08-20 11:08 AM - Giovanni Manghi

The problem is, that commands I have tried appear to be solved in a similar time using Windows and Linux. In other words, so far I have no evidence that GDAL/ogrprovider is the cause of the slowness. It appeared like the cause first, but then I realized that the command I tried is slow on Linux too. Sorry for the confusion.

I am confused now: is ogr slow or not with your dataset? Only on Windows (as the subject of this ticket suggest) or also on Linux?

#13 - 2018-08-20 12:42 PM - David Fiedler

Giovanni Manghi wrote:

The problem is, that commands I have tried appear to be solved in a similar time using Windows and Linux. In other words, so far I have no evidence that GDAL/ogrprovider is the cause of the slowness. It appeared like the cause first, but then I realized that the command I tried is slow on Linux too. Sorry for the confusion.

I am confused now: is ogr slow or not with your dataset? Only on Windows (as the subject of this ticket suggest) or also on Linux?

It is slow or quick depending on the command used. On both platforms. I think I'm lacking the insight into the ogr needed to evaluate it in a meaningful way.

#14 - 2018-08-20 05:57 PM - Giovanni Manghi

David Fiedler wrote:

Giovanni Manghi wrote:

The problem is, that commands I have tried appear to be solved in a similar time using Windows and Linux. In other words, so far I have no evidence that GDAL/ogrprovider is the cause of the slowness. It appeared like the cause first, but then I realized that the command I tried is slow on Linux too. Sorry for the confusion.

I am confused now: is ogr slow or not with your dataset? Only on Windows (as the subject of this ticket suggest) or also on Linux?

It is slow or quick depending on the command used. On both platforms.

can you make practical/real examples and possibly also add a download link for such a big gpx file? Thanks!

#15 - 2018-08-20 07:46 PM - David Fiedler

The file contains vehicle traces. I cannot send you the data, as it is private, but I can generate similar data for you, will it be helpful?

#16 - 2018-08-21 11:39 AM - Giovanni Manghi

David Fiedler wrote:

but I can generate similar data for you, will it be helpful?

yes

#17 - 2018-08-21 12:48 PM - David Fiedler

Giovanni Manghi wrote:

David Fiedler wrote:

but I can generate similar data for you, will it be helpful?

yes

Ok, I will provide you with the data. But it will take some time before I can get to it.

#18 - 2018-08-27 02:41 PM - David Fiedler

Here is the link to the generated traces: <https://mega.nz/#F!5fIXElqT!SZd6NISXoIQ-nXwkLqAJrQ>

For me, the loading time (for the largest file) on Windows is 11 minutes before feature selection and another 12 minutes until the rendering begins - 23 minutes of completely unresponsive UI in total.

#19 - 2018-08-30 05:22 PM - Giovanni Manghi

David Fiedler wrote:

Here is the link to the generated traces: <https://mega.nz/#F!5fIXElqT!SZd6NISXoIQ-nXwkLqAJrQ>

For me, the loading time (for the largest file) on Windows is 11 minutes before feature selection and another 12 minutes until the rendering begins - 23 minutes of completely unresponsive UI in total.

could you provide me with a direct http link, something I can get with wget? the above service does not work on Firefox (and I don't want to install chrome).

#20 - 2018-08-31 04:58 PM - David Fiedler

Giovanni Manghi wrote:

David Fiedler wrote:

Here is the link to the generated traces: <https://mega.nz/#F!5fIXElqT!SZd6NISXoIQ-nXwkLqAJrQ>

For me, the loading time (for the largest file) on Windows is 11 minutes before feature selection and another 12 minutes until the rendering begins - 23 minutes of completely unresponsive UI in total.

could you provide me with a direct http link, something I can get with wget? the above service does not work on Firefox (and I don't want to install chrome).

Here is a link to google drive: <https://drive.google.com/drive/folders/16GF1WVC1WfcLFKqZwuaBjSlpKZC6iU4t?usp=sharing>

I'm sorry, I wasn't aware that Mega does not support Firefox for large downloads.

#21 - 2018-08-31 05:10 PM - Giovanni Manghi

Here is a link to google drive: <https://drive.google.com/drive/folders/16GF1WVC1WfcLFKqZwuaBjSlpKZC6iU4t?usp=sharing>

this seems to point only to the smaller sample.

#22 - 2018-08-31 09:14 PM - David Fiedler

Giovanni Manghi wrote:

Here is a link to google drive: <https://drive.google.com/drive/folders/16GF1WVC1WfcLFKqZwuaBjSlpKZC6iU4t?usp=sharing>

this seems to point only to the smaller sample.

The link points to a folder with three files of different sizes. It pointed only to one file first, but I updated it within a minute, you had to click the link very fast - try it now.

#23 - 2018-09-02 09:04 PM - Giovanni Manghi

it is slow (but acceptable to me) on linux, but is ogr that is slow and I don't think this is a QGIS problem. To mt this should be closed as is an upstream issue.

```
giovanni@sibirica:~/Downloads$ time ogrinfo -so dummy_traces100000.gpx tracks
```

Had to open data source read-only.

```
INFO: Open of `dummy_traces100000.gpx'  
      using driver `GPX' successful.
```

Layer name: tracks

Geometry: Multi Line String

Feature Count: 100000

Extent: (14.257992, 49.942957) - (14.694051, 50.174524)

Layer SRS WKT:

```
GEOGCS["WGS 84",  
  DATUM["WGS_1984",  
    SPHEROID["WGS 84",6378137,298.257223563,  
      AUTHORITY["EPSG","7030"]],  
    AUTHORITY["EPSG","6326"]],  
  PRIMEM["Greenwich",0,  
    AUTHORITY["EPSG","8901"]],  
  UNIT["degree",0.01745329251994328,  
    AUTHORITY["EPSG","9122"]],  
  AUTHORITY["EPSG","4326"]]
```

name: String (0.0)

cmt: String (0.0)

desc: String (0.0)

src: String (0.0)

link1_href: String (0.0)

link1_text: String (0.0)

link1_type: String (0.0)

link2_href: String (0.0)

link2_text: String (0.0)

link2_type: String (0.0)

number: Integer (0.0)

type: String (0.0)

real 2m7,304s

user 2m5,906s

sys 0m1,364s

#24 - 2018-09-17 11:09 AM - David Fiedler

Giovanni Manghi wrote:

it is slow (but acceptable to me) on linux, but is ogr that is slow and I don't think this is a QGIS problem. To mt this should be closed as is an upstream issue.

```
giovanni@sibirica:~/Downloads$ time ogrinfo -so dummy_traces100000.gpx tracks
```

Had to open data source read-only.

INFO: Open of `dummy_traces100000.gpx`

using driver `GPX` successful.

Layer name: tracks

Geometry: Multi Line String

Feature Count: 100000

Extent: (14.257992, 49.942957) - (14.694051, 50.174524)

Layer SRS WKT:

GEOGCS["WGS 84",

DATUM["WGS_1984",

SPHEROID["WGS 84",6378137,298.257223563,

AUTHORITY["EPSG","7030"]],

AUTHORITY["EPSG","6326"]],

PRIMEM["Greenwich",0,

AUTHORITY["EPSG","8901"]],

UNIT["degree",0.01745329251994328,

AUTHORITY["EPSG","9122"]],

AUTHORITY["EPSG","4326"]]

name: String (0.0)

cmt: String (0.0)

desc: String (0.0)

src: String (0.0)

link1_href: String (0.0)

link1_text: String (0.0)

link1_type: String (0.0)

link2_href: String (0.0)

link2_text: String (0.0)

link2_type: String (0.0)

number: Integer (0.0)

type: String (0.0)

real 2m7,304s

user 2m5,906s

sys 0m1,364s

I don't think this explains the twenty minutes loading time of the file on Windows.

#25 - 2019-02-23 08:36 PM - Jürgen Fischer

- Status changed from Feedback to Closed

- Resolution set to no timely feedback

Bulk closing 82 tickets in feedback state for more than 90 days affecting an old version. Feel free to reopen if it still applies to a current version and you have more information that clarify the issue.