

QGIS Application - Bug report #1978

QGIS take ages to open big (or just complex) vectors

2009-10-05 09:00 AM - Giovanni Manghi

Status: Closed	
Priority: High	
Assignee:	
Category: Vectors	
Affected QGIS version: master	Regression?: No
Operating System: All	Easy fix?: No
Pull Request or Patch supplied: No	Resolution: fixed
Crashes QGIS or corrupts data: No	Copied to github as #: 12038
Description	
<p>Hi,</p> <p>I noticed that while trying open a quite big shapefile</p> <p>http://www.iucnredlist.org/resources/spatial/mammals/cetartiodactyla_v1223400057.zip</p> <p>Qgis freezes, or at least do not give sign of life for a long time. The same shapefile on other software (ex. gvSIG) opens in a matter of seconds.</p> <p>Using QGIS trunk (compiled today) under Ubuntu 9.04</p> <p>The same behaviour was reported by others with other big shapefiles.</p>	

History

#1 - 2009-11-25 11:04 AM - Giovanni Manghi

The link to that shapefile is no more available, use this instead

http://www.iucnredlist.org/spatial-data/MAMMALS_TERRESTRIAL.zip

Actually the behaviour changed as QGIS seems not to freeze anymore, it just take ages to open the shapefile and take even more to do any kind of basic (pan, zoom, etc.) operation. Using render caching do not seems to have any positive effect.

Also configuring QGIS to draw "x" features before updating the display (options -> rendering & svg) does not work for me.

Tested on both linux and windows xp with QGIS trunk.

#2 - 2009-11-25 11:10 AM - Giovanni Manghi

well, after all the render caching seems to have some kind of positive effect, but the time needed to go around in such a shapefile is still way too high (it takes just one minute instead of two).

I noticed also that is tremendously slow while opening the table of attributes.

#3 - 2009-11-30 01:34 AM - Giovanni Manghi

|

| Also configuring QGIS to draw "x" features before updating the display (options -> rendering & svg) does not work for me.

as confirmed in #2038, does not work if cache rendering is enabled.

#4 - 2009-12-07 04:18 AM - Giovanni Manghi

Well,

after a mail exchange with one of the main developers it results that this is not a QGIS issue but a QT issue, and so little can be done for now.

We can leave this open as reminder or close it. Opinions?

#5 - 2009-12-12 05:28 AM - Giovanni Manghi

It seems that with the fix for #2270 there was an improvement in the rendering performances of the big shapefile linked in this ticket (650mb of size). Still slow but now it is at least usable.

I still see serious speed issues using (for example) the identify/select feature tools, but zooming seems improved.

I especially noticed that if the identify tool results window is docked, then just selecting the identify tool make the canvas to be refreshed, hurting the user experience.

#6 - 2011-02-18 07:29 AM - Paolo Cavallini

The problem is still there. In real work with complex data this is really a stumbling block.

I checked, and on the same machine, [[ArcMap]] opens the same file much faster.

#7 - 2011-10-05 01:10 PM - Giovanni Manghi

- Target version changed from Version 1.7.0 to Version 2.0.0

- Assignee deleted (nobody -)

- Priority changed from Low to High

- Pull Request or Patch supplied set to No

#8 - 2011-11-09 09:55 PM - Nathan Woodrow

I suspect, and I might be wrong, that some of the reason that it's slow to open large ogr based layers is because QGIS doesn't limit the dataset on first load. It seems by looking at the code and the debug output that QGIS will just try and load the whole thing vs just loading the section from the file it needs. I'm guessing if we just loaded what we need + a little bit on each side we could increase loading performance quite a bit.

#9 - 2011-11-10 01:45 PM - Thaddeus -

- File TestLineString.py added

Load on demand: sounds like the best optimization to me!

I get an eight-time slower performance factor when activating the "Make lines appear less jagged at the expense of some drawing performance". Is this because this is a single core PC ? Have not tried on a multicore one yet.

Off <-- Make lines appear less jagged at the expense of some drawing performance.

1.7.1

- > load & display -> 7; 7; 6.4; 6.40; 6.40
- > zoom in to rectangle -> 6; 7; 6; 5.25; 5.29
- > zoom in to layer extents -> 6.44; 6.25

On <-- Make lines appear less jagged at the expense of some drawing performance.

1.7.1

- > load & display -> 39; 40; 48; 47; 47.81; 47.35
- > zoom in to rectangle -> 25; 27; 20
- > zoom in to layer extents -> 41; 44; 47; 47; 47.50

SAGA

- > load & display -> 8; 8.5
- > zoom in to rectangle -> 0.5
- > zoom in to layer extents -> 6; 6

Numbers in seconds; multiple trials; etc; etc;

OTF CRS disabled.

"Fix problems with incorrect filled polygons" always disabled.

Not using cache -- not repeating screens.

Wiping out ram: no Hardware optimization.

Processing times repeatable: consistent on several trials.

Always the same PC: P4, 2.8GHz, 3GBytes RAM, 1 core.

TestLineString.py

TestLineString.shp; 13MBytes

TestLineString.dbf; 1.2MBytes

Test100000.shx

#10 - 2011-11-11 02:54 AM - Giovanni Manghi

- *Must fix deleted (No)*

- *Category changed from Data Provider to Vectors*

Here is an interesting vector

http://www.faunalia.pt/downloads/test_slow_vector.zip

it is a (multi)polygon one, just 17 MB size, with many (complex) geometries and a very few attributes as they are aggregated in just 12 multipart polygons.

In its original state the vector is **very** slow (considered the size), especially when compared with other software, adding a spatial index and/or import in postgis doesn't change a lot. Toggling editing and try use the node tool freezes QGIS.

After converting the layer to a simplepart one it become much more easy to handle.

#11 - 2011-11-11 02:57 AM - Paolo Cavallini

- *Must fix set to yes*

#12 - 2011-11-11 04:15 AM - Giovanni Manghi

- *Subject changed from QGIS take ages to open big shapefiles to QGIS take ages to open big (or just complex) vectors*

#13 - 2011-11-11 06:30 AM - Martin Dobias

Thaddeus - wrote:

Load on demand: sounds like the best optimization to me!

I get an eight-time slower performance factor when activating the "Make lines appear less jagged at the expense of some drawing performance". Is this because this is a single core PC ? Have not tried on a multicore one yet.

We do load data on demand. The reason why zoomed in rendering took the same time to render is probably that you did not build spatial index for the shapefile.

The slow performance with antialiasing turned on is known, but there's hardly anything you can do against it - you need to choose whether you want nice or fast maps.

Multicore rendering is currently not supported. When it will be ready, it will primarily split rendering of various layers to cores - not splitting rendering of one layer to multiple cores.

#14 - 2011-11-11 06:43 AM - Martin Dobias

Giovanni Manghi wrote:

The link to that shapefile is no more available, use this instead

http://www.iucnredlist.org/spatial-data/MAMMALS_TERRESTRIAL.zip

Actually the behaviour changed as QGis seems not to freeze anymore, it just take ages to open the shapefile and take even more to do any kind of basic (pan, zoom, etc.) operation. Using render caching do not seems to have any positive effect.

Also configuring QGis to draw "x" features before updating the display (options -> rendering & svg) does not work for me.

Tested on both linux and windows xp with QGis trunk.

That's a nasty layer :-) Really, the shape data file itself has a size of 500MB, the polygons consist of lots of vertices, the polygons overlap a lot.

A lot of time is spent in OGR while reading (particularly in organizePolygons() method). On my box rendering takes ~20 seconds. How much time other software needs to render it?

Btw. what render caching does is storing the rendered image for a particular extent. So it helps only when you have other layers and modify other layers' appearance: on refresh the layers which did not change are drawn from the cache.

#15 - 2011-11-11 06:47 AM - Martin Dobias

Giovanni Manghi wrote:

Here is an interesting vector

http://www.faunalia.pt/downloads/test_slow_vector.zip

*it is a (multi)polygon one, just 17 MB size, with many (complex) geometries and a very few attributes as they are aggregated in just 12 multipart polygons. In its original state the vector is **very** slow (considered the size), especially when compared with other software, adding a spatial index and/or import in postgis doesn't change a lot. Toggling editing and try use the node tool freezes QGIS.*

After converting the layer to a simplepart one it become much more easy to handle.

This looks like a completely different problem, a separate ticket would be appropriate.

#16 - 2011-11-13 02:40 AM - Giovanni Manghi

This looks like a completely different problem, a separate ticket would be appropriate.

done #4523

#17 - 2011-11-13 02:59 AM - Giovanni Manghi

That's a nasty layer :-). Really, the shape data file itself has a size of 500MB, the polygons consist of lots of vertices, the polygons overlap a lot.

A lot of time is spent in OGR while reading (particularly in organizePolygons() method). On my box rendering takes ~20 seconds. How much time other software needs to render it?

Yes you are right, is a nasty shape, kind of a special case. I tested other software, gvSIG and saga, and the rendering time is not that much better compared with QGIS (more or less half the time), but probably the most interesting thing is that is possible to zoom while the geometries are rendered, so the user experience is quite good.

I would like to underline that since opening this ticket the performance of QGIS with such big vectors is much better, but it seems that there room for some improvement.

#18 - 2011-11-17 05:15 AM - Giovanni Manghi

Martin Dobias wrote:

Giovanni Manghi wrote:

Here is an interesting vector

http://www.faunalia.pt/downloads/test_slow_vector.zip

*it is a (multi)polygon one, just 17 MB size, with many (complex) geometries and a very few attributes as they are aggregated in just 12 multipart polygons. In its original state the vector is **very** slow (considered the size), especially when compared with other software, adding a spatial index and/or import in postgis doesn't change a lot. Toggling editing and try use the node tool freezes QGIS.*

After converting the layer to a simplepart one it become much more easy to handle.

This looks like a completely different problem, a separate ticket would be appropriate.

Wow, this has been really a roller-coaster lately. Rendering time of this vector seems **much** better now in qgis-master (as also selecting features) on Linux. Also the rendering time of the other huge vector linked in this ticket is **much** better (just above 10 seconds on a i3 machine), especially when the "rendering quality" options are not checked (acceptable to me especially for huge vectors).

Still not possible to edit them (problem assessed in #4523) as qgis freezes.

On my part this can be closed, given that performances will not drop again with the upcoming commits.

Can anyone check please with your own problematic/complex/big vectors?

#19 - 2011-11-18 12:26 AM - Giovanni Manghi

- Status changed from Open to Feedback

#20 - 2012-01-18 12:10 PM - Giovanni Manghi

- Crashes QGIS or corrupts data set to No

- Status changed from Feedback to Closed

- Resolution set to fixed

- Affected QGIS version set to master

To me this can be closed. Rendering time of big vectors is good now.

#21 - 2012-02-26 09:48 AM - Jorge Rocha

I'm having serious performance problems with the new 1.7.4 in Linux.

My test set (the administrative limits of PT) is available to download from

http://www.igeo.pt/produtos/cadastro/caop/download/CAOP2011_Shapes/Cont_AAAd_CAOP2011.zip

The shapefile or the same +- 4k polygons in PostGIS take almost the same time to load, which is more than one minute (about 1'20"). In 1.7.3, the same data used to load in less than 5 sec.

It's a huge performance difference. These low values are measured with default options, after installing a a fresh 1.7.4 QGIS.

As anybody has already experiencing such difference in 1.7.4?

```
$ apt-cache policy qgis
```

```
qgis:
```

```
Instalado: 1.7.4-1~oneiric1
```

```
Candidato: 1.7.4-1~oneiric1
```

```
Tabela de Versão: *** 1.7.4-1~oneiric1 0
```

```
500 http://ppa.launchpad.net/ubuntuqgis/ubuntuqgis-unstable/ubuntu/ oneiric/main i386 Packages
```

#22 - 2012-02-26 09:53 AM - Giovanni Manghi

Olá Jorge,

are you sure you have the exactly same configurations you had in 1.7.3? I know well your dataset and I will try soon 1.7.4 under Linux so I will leave more feedback later.

#23 - 2012-02-26 10:29 AM - Jorge Rocha

Gio,

I was able to reboot with my former disk, with 1.7.3 QGIS. It takes 2 sec. to open the shapefile in 1.7.3 , against about 1'20" with 1.7.4.

```
$ apt-cache policy qgis
```

```
qgis:
```

```
Instalado: 1.8.0+git20110612+a954488~maverick1
```

```
Candidato: 1.8.0+git20110612+a954488lem.~maverick1
```

```
Tabela de Versão: *** 1.8.0+git20110612+a954488~maverick1 0
```

```
100 /var/lib/dpkg/status
```

```
1.7.3-1~oneiric1 0
```

```
500 http://ppa.launchpad.net/ubuntuqgis/ubuntuqgis-unstable/ubuntu/ oneiric/main amd64 Packages
```

```
libqtcore4 4:4.7.4-0ubuntu8.1 Qt 4 core module
```

```
libgeos-3.2.2
```

```
libgdal1-1.8.0
```

```
3.0.0-15-generic x86_64
```

#24 - 2012-02-26 03:22 PM - Jorge Rocha

Jorge Rocha wrote:

I'm having serious performance problems with the new 1.7.4 in Linux.

My test set (the administrative limits of PT) is available to download from

http://www.igeo.pt/produtos/cadastro/caop/download/CAOP2011_Shapes/Cont_AAAd_CAOP2011.zip

The shapefile or the same +- 4k polygons in PostGIS take almost the same time to load, which is more than one minute (about 1'20"). In 1.7.3, the same data used to load in less than 5 sec.

It's a huge performance difference. These low values are measured with default options, after installing a a fresh 1.7.4 QGIS.

As anybody has already experiencing such difference in 1.7.4?

```
$ apt-cache policy qgis
```

```
qgis:
```

```
Instalado: 1.7.4-1~oneiric1
```

```
Candidato: 1.7.4-1~oneiric1
```

```
Tabela de Versão: - 1.7.4-1~oneiric1 0
```

```
500 http://ppa.launchpad.net/ubuntuqgis/ubuntuqgis-unstable/ubuntu/ oneiric/main i386 Packages
```

```
libqtcore4 4:4.7.4-0ubuntu8.1
```

```
libgeos-3.2.2
```

```
libgdal1-1.8.0
```

```
3.0.0-16-generic-pae i686
```

libqt, libgeos and libgdal are the same, testing using 1.7.3 and 1.7.4. The performance is completely different.

#25 - 2012-02-27 11:37 AM - Giovanni Manghi

Olá Jorge

| *libqt, ibgeos and libgdal are the same, testing using 1.7.3 and 1.7.4. The performance is completely different.*

I made a very clean install of QGIS 1.7.4 (from standard repo) in my 6 years old laptop (Ubuntu 11.10 64 bit), and it seems I cannot confirm your issue. The vector opens in a few instants with the "make lines appear less jagged..." option OFF and a few seconds (4/5) with the option ON.

It takes around 30 seconds to open this huge http://www.gadm.org/data2/gadm_v2_shp.zip shapefile.

#26 - 2012-02-27 03:45 PM - Jorge Rocha

HOW TO REPRODUCE THE BUG

--

in Ubuntu 11.10 with QGIS 1.7.4

--

This bug can be reproduced simply by removing (or moving) the existing preferences. In the absence of any preferences, the rendering is really slow.

1) `mv ~/.config/QuantumGIS/QGIS.conf ~/.config/QuantumGIS/QGIS.conf.old`

2) `qgis --lang en --nologo --noplugins`

2.1) opening the mentioned shapefile will take a lot of time (> 1 minute)

AVOID THE BUG

--

This bug can be avoided simply by tweaking some preferences, to ensure the existence of the preferences file.

2.2) change some preferences, in Edit -> Options

2.3) Close qgis

3) start qgis after tweaking some preferences: the good performance is back.

BETTER PERFORMANCE

--

As usual, better rendering performance can be archived by tweaking the Render quality options in the Render tab, in Edit -> Options.

#27 - 2012-02-28 01:53 AM - Giovanni Manghi

Olá Jorge,

that sounds odd, and if confirmed it will deserve a separate ticket. Personally I still cannot confirm, even after deleting `.config/QuantumGIS` and `.qgis`

Can you test on other machines?

abraços

Jorge Rocha wrote:

HOW TO REPRODUCE THE BUG

2024-04-27

8/9

--
in Ubuntu 11.10 with QGIS 1.7.4
--
This bug can be reproduced simply by removing (or moving) the existing preferences. In the absence of any preferences, the rendering is really slow.

1) `mv ~/.config/QuantumGIS/QGIS.conf ~/.config/QuantumGIS/QGIS.conf.old`

2) `qgis --lang en --nologo --noplugins`

2.1) opening the mentioned shapefile will take a lot of time (> 1 minute)

AVOID THE BUG

--
This bug can be avoided simply by tweaking some preferences, to ensure the existence of the preferences file.

2.2) change some preferences, in Edit -> Options

2.3) Close qgis

3) start qgis after tweaking some preferences: the good performance is back.

BETTER PERFORMANCE

--
As usual, better rendering performance can be archived by tweaking the Render quality options in the Render tab, in Edit -> Options.

#28 - 2012-03-02 03:27 AM - Jorge Rocha

This bug occurs in every new installed Ubuntu, using 1.7.4 without any previous QGIS installation.

#29 - 2012-03-02 03:30 AM - Giovanni Manghi

Jorge Rocha wrote:

This bug occurs in every new installed Ubuntu, using 1.7.4 without any previous QGIS installation.

Olá Jorge, please file a new ticket. Abraços.

Files

TestLineString.py	859 Bytes	2011-11-10	Thaddeus -
-------------------	-----------	------------	------------