

QGIS Application - Bug report #8101

"Save as" on a large PostGIS point layer doesn't work

2013-06-18 09:26 AM - Aren Cambre

Status:	Closed	
Priority:	Low	
Assignee:		
Category:	Vectors	
Affected QGIS version:	master	Regression?: No
Operating System:		Easy fix?: No
Pull Request or Patch supplied:		Resolution: end of life
Crashes QGIS or corrupts data:		Copied to github as #: 16936
Description		
<p>On 1.8.0 and 1.9 192e130, I load a PostGIS layer with 2,889,598 points. These represent a few years' worth of a type of event along roadways in a state of the United States. I right-click on the layer, select Save As, select a path and filename, check Skip attribute creation (the layer has a lot of attributes; don't need them in the SHP), and press OK. I don't change the Format field from ERSI Shapefile.</p> <p>QGIS just hangs for a few minutes. It eventually gives the attached error. It does produce a SHP (also attached in a ZIP), but the file sizes are far too small. The DBF only has 1253 data rows.</p> <p>(Just to be clear, I can't share the data due to state law and confidentiality agreements.)</p>		

History

#1 - 2013-06-18 11:25 AM - Giovanni Manghi

- Status changed from Open to Feedback

I tested with a 568k polygon layer on a remote postgis server, and it took about 2 minutes to export it as shape (without attributes).

I will see if I can find something bigger...

#2 - 2013-06-18 11:48 AM - Aren Cambre

This is interesting. I get the same symptom if I use the **Filter** feature to bring the layer down to about 6000 points.

Checking on other factors...

#3 - 2013-06-18 12:09 PM - Aren Cambre

OK, on the export of the filtered subset of only about 6000 points, it takes about 6 minutes to complete. I've attached the output of the error dialog.

While this is happening, QGIS is using 100% of one of my CPUs.

Postgres has little activity, verified both through CPU usage and pgAdmin III's Server Status feature. All it does is receive occasional **FETCH FORWARD 2000 FROM qgisf0** commands every minute or two, and that appears to produce 12 locks.

Something that could be distinctive is this PostGIS table has 148 columns (including the numeric index and geometry columns).

#4 - 2013-06-18 12:09 PM - Aren Cambre

- File *heatmap_errors.txt* added

Here's that file of the heatmap errors.

#5 - 2013-06-18 12:19 PM - Aren Cambre

Interesting: I ran the same operation back to back (with the filtered dataset). It took 6 minutes the first time, 19 minutes the second time.

#6 - 2013-06-18 12:51 PM - Aren Cambre

Just tried again but without checking the **Skip attribute creation** box. The only real difference here was that the DBF part of the ShapeFile was huge, as expected, due to all the attributes. It still took a very long time.

#7 - 2013-07-05 06:18 PM - Aren Cambre

I have a SQL dump of the PostGIS datasource handy, so I can share it with you. Please email me at aren@arencambre.com to get it.

#8 - 2013-07-09 04:54 AM - Giovanni Manghi

- Status changed from *Feedback* to *Open*
- Priority changed from *Normal* to *Low*
- Category changed from *GUI* to *Vectors*
- Subject changed from "Save as" on large Postgres point layer doesn't work to "Save as" on large a PostGIS point layer doesn't work

This are my findings:

- exporting large/huge postgis layers works ok: I tested on layers with up to 5 million features and in general it works ok (with "save as...")
- I confirm that with **your** dataset the operation (save as...) fails after a few minutes with a list of

Feature creation error (OGR error: Pointer 'hFeat' is NULL in 'OGR_L_SetFeature').

)Stopping after 1001 errors

errors

- exporting the same dataset with ogr2ogr is ok and gives the expected shapefile
- exporting with pgsq2shp seems to work fine, but beside not returning any error the number of features is not the same as the original postgis layer

#9 - 2013-07-09 05:12 AM - Matthias Kuhn

There's a hardcoded limit of 1000 allowed errors in QgsVectorFileWriter (qgsvectorfilewriter.cpp:860).

This could (and should) be made configurable. But this leaves open the question of which features fail to be written and for what reason.

#10 - 2013-07-09 05:14 AM - Nathan Woodrow

Why is there a hard coded limit? That is pretty bad design IMO. We should just write what we can and report what we couldn't at the end.

#11 - 2013-07-09 05:15 AM - Giovanni Manghi

- Subject changed from "Save as" on large a PostGIS point layer doesn't work to "Save as" on a large PostGIS point layer doesn't work

Matthias Kuhn wrote:

There's a hardcoded limit of 1000 allowed errors in QgsVectorFileWriter (qgsvectorfilewriter.cpp:860).
This could (and should) be made configurable. But this leaves open the question of which features fail to be written and for what reason.

as I said the operation in general works ok even for very large postgis vectors. With this particular dataset there is this problem.

#12 - 2013-07-09 06:46 AM - Aren Cambre

Is there anything unique about this dataset? I haven't noticed any problems viewing it. It's just a bunch of points.

#13 - 2014-05-21 07:58 AM - Bill Morris

I've just run into the same error on a point file with 100 features and a single attribute column. Only 86 out of 100 exported. This bug is still out there.

<https://gist.github.com/wboykinm/751a02d51f95ddb3f5cd>

Aren Cambre wrote:

Is there anything unique about this dataset? I haven't noticed any problems viewing it. It's just a bunch of points.

#14 - 2017-05-01 01:09 AM - Giovanni Manghi

- Easy fix? set to No
- Regression? set to No

#15 - 2019-03-09 04:04 PM - Giovanni Manghi

- Resolution set to end of life
- Status changed from Open to Closed

End of life notice: QGIS 2.18 LTR

Source:

<http://blog.qgis.org/2019/03/09/end-of-life-notice-qgis-2-18-ltr/>

Files

save_as_error.png	28.1 KB	2013-06-18	Aren Cambre
test.zip	28.8 KB	2013-06-18	Aren Cambre
heatmap_errors.txt	46.9 KB	2013-06-18	Aren Cambre