

QGIS Application - Bug report #12630

Get sample query expression non-responsive for large PostGIS datasets

2015-04-23 03:07 AM - Pete Lelliott

Status:	Closed	
Priority:	Severe/Regression	
Assignee:		
Category:	Data Provider/PostGIS	
Affected QGIS version:	master	Regression?: No
Operating System:		Easy fix?: No
Pull Request or Patch supplied:	No	Resolution: fixed/implemented
Crashes QGIS or corrupts data:	No	Copied to github as #: 20745
Description		
<p>Apologies if I've put this into the wrong category, this is my first time posting a bug report.</p> <p>I have a large dataset with more than 3 billion points. When constructing a sql expression in the Set Filter of layer properties and using the 'Sample' feature, QGIS 2.8 goes non-responsive. The field I'm sampling is of type integer.</p> <p>Upon looking at the PostGIS server status, I can see that the query sent by QGIS is the following:</p> <pre>SELECT DISTINCT "field_to_sample"::text FROM "public"."feature_dataset" ORDER BY "field_to_sample" LIMIT 25</pre> <p>Within PostGIS, changing this to the actual datatype of the dataset returns a result in a little under 2 seconds;</p> <pre>SELECT DISTINCT "field_to_sample"::integer FROM "public"."feature_dataset" ORDER BY "field_to_sample" LIMIT 25</pre> <p>I imagine that the type conversion to text of large datasets (if that is what is happening?) could be causing the non-responsive behaviour within QGIS. Does the type conversion have to happen within the query?</p> <p>QGIS Version 2.6 does not do this type conversion, so therefore returns the expected result within a few seconds.</p> <p>Cheers,</p> <p>Pete</p>		

Associated revisions

Revision e509fc53 - 2015-05-10 10:55 AM - Jürgen Fischer

postgres provider: cast result of a subquery in min/max/uniqueValue(s) (fixes #12630; followup bf56457 ff)

Revision a9cccd2c - 2015-05-28 12:11 AM - Jürgen Fischer

postgres provider: cast result of a subquery in min/max/uniqueValue(s) (fixes #12630; followup bf56457 ff)

(cherry picked from commit e509fc53ce5cfdb64b984177c5279a6cf7d8644c)

History

#1 - 2015-04-26 03:15 PM - Giovanni Manghi

- Target version set to Version 2.8.2
- Category changed from GUI to Data Provider/PostGIS
- Priority changed from Normal to Severe/Regression

I don't have such big dataset to test with right now, anyway until then I'll tag this a regression as it seems to be.

#2 - 2015-04-27 04:07 AM - Giovanni Manghi

- OS version deleted (2.8 64 Bit)
- Affected QGIS version changed from 2.8.1 to master
- Operating System deleted (Windows 7 Professional 64Bit)

I confirmed on 2.6.1 on a large table (2.8M records): there is considerably faster and int columns values are ordered as numbers, not as text as in 2.8.1 and master.

#3 - 2015-05-10 01:56 AM - Jürgen Fischer

- Status changed from Open to Closed

Fixed in changeset commit:"e509fc53ce5cfdb64b984177c5279a6cf7d8644c".

#4 - 2015-05-10 03:51 AM - Jürgen Fischer

- Status changed from Closed to In Progress

#5 - 2015-05-10 03:52 AM - Jürgen Fischer

- Status changed from In Progress to Closed
- Resolution set to fixed/implemented