

## QGIS Application - Bug report #6232

### Postgis add layer dialog: Cancel slow operations

2012-08-21 12:27 AM - Matthias Kuhn

<b>Status:</b>	Closed	
<b>Priority:</b>	Normal	
<b>Assignee:</b>		
<b>Category:</b>	Data Provider/PostGIS	
<b>Affected QGIS version:</b>	master	<b>Regression?:</b> No
<b>Operating System:</b>		<b>Easy fix?:</b> No
<b>Pull Request or Patch supplied:</b>		<b>Resolution:</b> fixed
<b>Crashes QGIS or corrupts data:</b>		<b>Copied to github as #:</b> 15541
<b>Description</b>		
<p>When I click on "Add postgis layer" and select a database and click "Connect", all the schemes appear pretty fast and I can select a table (in ~2s). To fully "load" the database (i.e. when the text of the button "Stop" changes to "Connect" and the button "Add" is enabled) it takes much longer (&gt;20s). I don't know, what information is loaded in the additional time.</p> <p>There are a number of problems with this associated here:</p> <p>Problems:</p> <ul style="list-style-type: none"><li>- Why am I not allowed to select a table, although QGIS already knows its name and primary key?</li><li>- If I click "Close" before the database is fully "loaded", it still waits for this time to pass, before allowing me any further action on the UI.</li></ul> <p>(Making the application appear slow and unresponsive)</p> <ul style="list-style-type: none"><li>- If I click "Close" before the database is fully "loaded", the cursor does not turn back from hourglass to pointer. This indicates to me that I should not continue my work, although you can continue to work after this timeout has passed.</li></ul> <p>Questions:</p> <ul style="list-style-type: none"><li>- What could it be, that the dialog is waiting for this long time and is it necessary to wait for this?</li><li>- If the user hits "Close" can this process be stopped immediately (let's say within max. 3s)?</li><li>- Can the ongoing process be forced to the background? (It's a thread that QgsPgSourceSelect waits for in its destructor.)</li><li>- If the above can not be accomplished, I think at least we should:<ul style="list-style-type: none"><li>- keep the select dialog open until the operation finished. So the user knows, what he is waiting for.</li><li>- reset the cursor to it's default state when the operation completed</li></ul></li></ul>		

### History

#1 - 2012-08-21 02:25 AM - Jürgen Fischer

| *What could it be, that the dialog is waiting for this long time and is it necessary to wait for this?*

If your geometry type and srids are not fully - either because it's not in geometry\_columns or not listed with a point, line or polygon type in geometry\_columns, QGIS needs to find which geometry types and srids are available.

| *If the user hits "Close" can this process be stopped immediately (let's say within max. 3s)?*

If you close the detection thread is notified to not run any more queries - but as the queries are blocking (that's why there is a thread) the thread can only react on that once the running query finishes.

| *Can the ongoing process be forced to the background? (It's a thread that QgsPgSourceSelect waits for in its destructor.)*

Not sure. That might into race conditions.

## #2 - 2012-08-21 04:33 AM - Matthias Kuhn

If your geometry type and srids are not fully - either because it's not in geometry\_columns or not listed with a point, line or polygon type in geometry\_columns, QGIS needs to find which geometry types and srids are available.

Yes, we have a couple of views here that have not yet been inserted into the geometry\_columns. As far as I can see, there is no way to add them automated ([http://postgis.org/documentation/manual-1.5/Probe\\_Geometry\\_Columns.html](http://postgis.org/documentation/manual-1.5/Probe_Geometry_Columns.html) states it does not scan for views)

So maybe in addition to this bug there should be a feature request, to

- a) only scan for geometry\_columns once a schema has been opened
  - b) ask the user if he really wants to scan for this (on bigger databases this can easily take a couple of minutes)
  - c) inform the user, how this can be improved and maybe even offer the possibility to add the missing entries.
- 

Concerning the cancellation of a request. I've looked at the PG docs and found the method PQcancel. I've never worked with this lib but this method looks promising to me and should be thread-safe. I think there is no reason to let the query complete, as the result will be discarded anyway?

## #3 - 2012-08-21 04:55 AM - Jürgen Fischer

Matthias Kuhn wrote:

*So maybe in addition to this bug there should be a feature request, to*  
*a) only scan for geometry\_columns once a schema has been opened*

That only helps if you have multiple schemas and even then a single schema might be enough to show the same behaviour.

*b) ask the user if he really wants to scan for this (on bigger databases this can easily take a couple of minutes)*

You can already setup your connection that way - and use check "use estimated metadata", which should help a lot, if you're dealing with large databases.

*c) inform the user, how this can be improved and maybe even offer the possibility to add the missing entries.*

In PostGIS 2.0 geometry\_columns isn't a table anymore anyway.

*Concerning the cancellation of a request. I've looked at the PG docs and found the method PQcancel. I've never worked with this lib but this method looks promising to me and should be thread-safe. I think there is no reason to let the query complete, as the result will be discarded anyway?*

Good point. I'll take a look at that.

## #4 - 2012-08-21 05:22 AM - Matthias Kuhn

Jürgen Fischer wrote:

*Matthias Kuhn wrote:*  
*That only helps if you have multiple schemas and even then a single schema might be enough to show the same behaviour.*

True. It's probably the wrong approach.

*You can already setup your connection that way - and use check "use estimated metadata", which should help a lot, if you're dealing with large database.*

Yes, I've just been told about this possibility. It's just not too intuitive. I wouldn't know if I didn't have somebody to tell me. I still think it would be a good idea to print some information next to the affected tables/views and tell the user, what the problem is, and what can be done about it. A yellow warning triangle and a popup with a short explanation would be good.

*In PostGIS 2.0 geometry\_columns isn't a table anymore anyway.*

So this problem is solved there by design? I wonder how many people are still using PG 1.x and when everybody will have updated?

*Good point. I'll take a look at that.*

Great

#### **#5 - 2012-08-21 05:45 AM - Jürgen Fischer**

Matthias Kuhn wrote:

*Good point. I'll take a look at that.*

*Great*

implemented in commit:b6760ce1d4

#### **#6 - 2012-08-22 05:33 AM - Matthias Kuhn**

- Status changed from Open to Closed
- Resolution set to fixed

Perfect. Works like a charm.

I'll mark this closed as the issue in the title is fixed.

#### **#7 - 2012-09-19 04:00 AM - hamish -**

Hi,

just as a point of info for anyone in the same situation as me:

I had to change the geometry\_columns type to MULTILINESTRING instead of LINE, otherwise it just got stuck on "Detecting..." forever (or at least >10 minutes, when I canceled it). POINTS was fine as POINTS, I haven't needed to try POLYGON vs. MULTIPOLYGON but perhaps there are similar issues there.

Thanks for the tip, connecting to the PostGIS db (and/or going into the Query Builder) now takes 2 seconds instead of 4 minutes for these large datasets, even without turning on the 'estimate table metadata' tick box in the server settings (glad to know about that trick too).

regards,  
Hamish

**#8 - 2012-09-19 05:34 AM - Jürgen Fischer**

hamish - wrote:

| *I had to change the geometry\_columns type to MULTILINESTRING instead of LINE,*

Was that originally LINE or LINESTRING?