

QGIS Application - Bug report #10619

Attribute table 'Show features visible on map' option does not speed up loading of large postgis table

2014-06-18 01:25 AM - Richard Duivenvoorde

Status:	Closed	
Priority:	High	
Assignee:	Jürgen Fischer	
Category:	Attribute table	
Affected QGIS version:	master	Regression?: No
Operating System:		Easy fix?: No
Pull Request or Patch supplied:		Resolution: worksforme
Crashes QGIS or corrupts data:		Copied to github as #: 19023
Description I have read in one of the postgis dump/backups from here: http://data.nlextract.nl/bag/postgis/ actually: http://data.nlextract.nl/bag/postgis/bag-8mrt2014.backup which has all addresses and buildings in the netherlands (>10 million records each table). Because opening the attribute table isn't possible for such table, Jurgen pointed me to the Data Sources settings, where I selected 'Show features visible on map' as Attribute table behaviour. Even zoomin in at one address or house, it still seems QGIS is loading ALL features before being able to just show me that one record. This takes >5 minutes, while painting those houses is second work (because of the spatial query). This settings option IS working with me on shapefiles though. Can somebody confirm? (to test you could read in that dump)		
Related issues:		
Related to QGIS Application - Bug report # 12318: attribute table does not sh...	Closed	2015-03-05
Related to QGIS Application - Feature request # 12926: Improve attribute tabl...	Closed	2015-06-09
Related to QGIS Application - Bug report # 13208: Poor attribute query and sp...	Closed	2015-08-11

Associated revisions

Revision 86d08b6c - 2015-02-14 09:45 PM - Jürgen Fischer

restore option to open the attribute table with only currently visible features
to improve performance (actually to make it usable with large tables again).

When used as initial setting in options the open attribute table will have a
fixed spatial filter to the canvas extent when it was opened. That extent is
also highlighted to indicate the filter.

fixes #10619

History

#1 - 2014-06-19 03:26 AM - Richard Duivenvoorde

update: confirmed here by myself on osgeo4w QGIS master (windows), connecting to another linux postgresql server.

it seems to work with smaller (postgres) tables and shapefiles though

#2 - 2014-06-19 04:17 AM - Giovanni Manghi

- Affected QGIS version changed from 2.2.0 to master

Richard Duivenvoorde wrote:

update: confirmed here by myself on osgeo4w QGIS master (windows), connecting to another linux postgresql server.

it seems to work with smaller (postgres) tables and shapefiles though

I can confirm, seems that 2.2 behave the same way.

#3 - 2014-06-20 12:07 AM - Matthias Kuhn

- Category changed from Data Provider/PostGIS to Attribute table

To be precise:

The option works, it just doesn't reveal the performance boost you expect?

#4 - 2014-06-20 12:38 AM - Richard Duivenvoorde

- File retrievingAllFeaturesButShowOnlyViewedSelected.png added

- File onlyviewedattr.png added

@matthias nope, it's not the performance, it's the fact that if I 'Show features visible on map' and I view only 10 features in the map (from this big dataset), that I'm confronted with the same in between dialog as if I select 'Show all features' (see onlyviewedattr.png) counting from 10000 to 15.000.000 in this case.

While the intended behaviour would be that the provider uses the extent of the current mapcanvas to do a spatial selection to retrieve that 10 records (instead of apparently trying to retrieve 15000000 features and then find the 10 which are in view currently).

But maybe I'm wrong. Maybe QGIS has not other way to first retrieve all features to show me the 10 which are in view.

But in that case (and in current status) I'm not able to open the attribute table with a postgis layer with that much records because whatever option I choose for 'Attribute table behaviour' QGIS takes 5 minutes to retrieve them and then does not show them to me anyway....

And if I abort the loading of the features, I'm not able to see the attributes of the current features in view, because those are apparently not loaded... (see retrievingAllFeaturesButShowOnlyViewedSelected.png)

#5 - 2014-06-20 02:31 AM - Giovanni Manghi

- Priority changed from High to Severe/Regression

surprise! in QGIS 1.8 it works like a charm (as expected).

#6 - 2014-06-20 02:42 AM - Giovanni Manghi

- Subject changed from 'Attribute table behaviour'-options not working with large postgis table to Attribute table 'Show features visible on map' option not

working with large postgis table

#7 - 2014-06-20 02:56 AM - Matthias Kuhn

- Priority changed from Severe/Regression to High

I wouldn't tag this as a blocker unless you want to postpone the release for a longer time.
That's not a quick fix.

QGIS 2.x does filtering differently and will e.g. adjust the visible items on the table when you zoom out.
QGIS 1.8 just loaded the features once on table init.

The "options" defined in the settings are also not constraints for the attribute table, but rather default values for the filtering mode you can choose on the attribute table.

#8 - 2014-06-20 03:01 AM - Giovanni Manghi

Matthias Kuhn wrote:

*I wouldn't tag this as a blocker unless you want to postpone the release for a longer time.
That's not a quick fix.*

I don't want to enter the discussion if a longer testing/fixing time has to be granted but among the list of actual regressions this seems like a huge one compared to others. I would suggest to leave the tag blocker even if no one will fix it on time for the release, as a reminder for the development of 2.6.

#9 - 2014-06-20 03:07 AM - Matthias Kuhn

I wouldn't say it's a huge one if noone complained in the +6 months the issue exists.

/me would vote for leaving it unfixed in order to get 2.4 out in reasonable time and to be free to implement a proper solution for 2.6.
YMMV

#10 - 2014-06-28 07:45 AM - Jürgen Fischer

- Target version changed from Version 2.4 to Future Release - High Priority

#11 - 2014-06-28 09:14 AM - Giovanni Manghi

*- Target version changed from Future Release - High Priority to Version 2.6
- Priority changed from High to Severe/Regression*

Now that 2.4 is out I'll re-tag this as regression.

#12 - 2014-06-28 09:57 AM - Jürgen Fischer

- Target version changed from Version 2.6 to Future Release - High Priority

#13 - 2014-07-22 03:17 PM - Matthias Kuhn

- Subject changed from Attribute table 'Show features visible on map' option not working with large postgis table to Attribute table 'Show features visible on map' option does not speed up loading of large postgis table

#14 - 2014-10-04 12:32 PM - Luigi Pirelli

Analysing with Jurgen and discussed with Giovanni Manghi we decided to postpone the solution to the next release for this reason:

- 1) is not really blocker (fetching optimization effectively works on rendering)
- 2) it doesn't work the optimization on fetching table view due the fact that if allow the "feature" to change dynamically among filter options
- 3) a combination of "main" visible option and "on table view" dynamic filtering can be implemented, but it touches the table model component that is used in other part of QGIS

So our evaluation is that it's better to apply the patch after next release and leaving time to test deeply in case of regressions.

#15 - 2014-10-06 03:04 AM - Matthias Kuhn

I agree that this should be postponed to after feature freeze.

I see two possibilities that can be evaluated to tackle this:

- Running the filter as a query on provider side every time the filter changes. This will lose the context (scroll state, selection, can probably be partially restored) of the table every time the filter changes (zoom, pan) and will need to re-populate it.
- Having a look at the main model's `canFetchMore()` and `fetchMore()` methods which allow to incrementally populate tables. With this approach, the table would not be fully populated initially, reducing initial loading time. Probably a feature iterator would need to be kept open to continue filling for subsequent scroll operations. I don't know exactly what the properties of such a solution are. Maybe we can have a threshold amount of features that will switch the table to this behavior.

#16 - 2014-10-06 04:31 AM - Richard Duivenvoorde

One other observation (not 100% sure it is related),

if I load this table in QGIS (zoomed in a lot to not load all 13.000.000 records), and select 10 houses and do a 'save as' and choose only selection (10 houses), QGIS also takes minutes to do this.

I think that is because of the problem.

Just want to add this as an argument to fix this (in a next version) as in my opinion you cannot use QGIS in this way on bigger postgres tables, if both loading the attributes, and saving a selection do not work

#17 - 2014-10-21 02:46 AM - Jürgen Fischer

- Assignee deleted (*Jürgen Fischer*)

#18 - 2014-10-21 09:10 PM - Martin Dobias

- Priority changed from *Severe/Regression* to *High*

Does not look like a blocker to me.

#19 - 2015-02-14 11:37 AM - Jürgen Fischer

- Assignee set to *Jürgen Fischer*

#20 - 2015-02-14 12:46 PM - Jürgen Fischer

- Status changed from *Open* to *Closed*

Fixed in changeset commit:"86d08b6c39bad37e4c9b468f55c4c60981b60fa2".

#21 - 2016-05-12 12:54 AM - Richard Duivenvoorde

- Status changed from Closed to Reopened

Hitting this again now.

Working with a Postgis db loaded with Openstreetmap data of the whole world.

Currently number of records for roads: 93.011.512

Trying to view records with attribute table. Zoomed in at streetlevel.

I can now 'Abort' the fetching (though the button seems unresponsive at times), so I can at least stop the fetching.

But the filtering is done on fetched records yes?

Would an option be to set a treshold at (user defined): 100000 records, And IF that number of features is received, stop and ask the user if he really wants to load all (maybe showing the total records, if provider has a count(*))?

OR maybe wants to add a spatial filter (for current extent of mapcanvas)?

Not sure if this is doable for all providers, but it would make working with this kind of datasets possible I think

#22 - 2016-05-24 03:10 AM - Giovanni Manghi

Richard Duivenvoorde wrote:

| *Hitting this again now.*

confirmed here too.

#23 - 2016-06-20 09:36 AM - Jürgen Fischer

- Status changed from Reopened to Feedback

Richard Duivenvoorde wrote:

| *Hitting this again now.*

And you have "Attribute table behaviour" set to "Show features visible on map" in options? Apparently that still works fine for me.

#24 - 2016-07-03 12:47 AM - Giovanni Manghi

Jürgen Fischer wrote:

Richard Duivenvoorde wrote:

Hitting this again now.

And you have "Attribute table behaviour" set to "Show features visible on map" in options? Apparently that still works fine for me.

I just tested on the latest master and it now works fine also here.

#25 - 2016-07-03 01:02 AM - Richard Duivenvoorde

will test tomorrow and report back

#26 - 2017-01-02 06:18 AM - Giovanni Manghi

- *Status changed from Feedback to Closed*

- *Resolution set to worksforme*

Closing for lack of feedback, please reopen if necessary.

Files

retrievingAllFeaturesButShowOnlyViewedSelected.png	151 KB	2014-06-19	Richard Duivenvoorde
onlyviewedattr.png	46.6 KB	2014-06-19	Richard Duivenvoorde