

QGIS Application - Bug report #20260

labelling is much slower in 3.4 compared to 2.18

2018-10-29 02:55 PM - Mario Baranzini

Status:	Closed	
Priority:	High	
Assignee:		
Category:	Labelling	
Affected QGIS version:	3.5(master)	Regression?: Yes
Operating System:		Easy fix?: No
Pull Request or Patch supplied:	No	Resolution:
Crashes QGIS or corrupts data:	No	Copied to github as #: 28081
Description		
<p>QGIS takes a very very long time (or stops) trying to show labels with placement around centroid for a layer containing polygons with a big number of parts.</p> <p>The problem is reproducible with the gadm36_0.shp (containing 256 polygons with a total of more than 100'000 parts) that can be downloaded from https://biogeo.ucdavis.edu/data/gadm3.6/gadm36_levels_shp.zip (the file is freely available for non-commercial use but redistribution is not allowed).</p> <p>The problem may depend on the geometry validity check. In fact, using the native processing algorithm with validity check activated, the same happens, but with the validity check disabled in the settings, there are no problems and centroids are calculated in a few seconds.</p> <p>See also: #20260-6</p>		

Associated revisions

Revision ff5a8bc0 - 2018-11-09 08:13 PM - Nyal Dawson

Revert "[pal] Use QgsGeometry::makeValid instead of buffer(0) to repair geometries"

This reverts commit e05a5a930241ec3c9c5df7880650da60382b956c.

The makeValid call is much slower than the previous "buffer(0)" approach

Fixes #20260

Revision 62d0bae6 - 2018-11-10 07:16 AM - Nyal Dawson

Revert "[pal] Use QgsGeometry::makeValid instead of buffer(0) to repair geometries"

This reverts commit e05a5a930241ec3c9c5df7880650da60382b956c.

The makeValid call is much slower than the previous "buffer(0)" approach

Fixes #20260

(cherry picked from commit ff5a8bc0aebf51cd34c82ed0feaac21cbd22e2ce)

History

#1 - 2018-10-29 09:57 PM - Giovanni Manghi

- Status changed from Open to Feedback
- Operating System deleted (Linux)

For a 1.2gb zipped shapefile I think is safe to say that if you don't limit labelling (by max number of labels, by scale, etc.) is safe to say that is expected to be slow.

#2 - 2018-10-30 06:55 AM - Mario Baranzini

Giovanni Manghi wrote:

For a 1.2gb zipped shapefile I think is safe to say that if you don't limit labelling (by max number of labels, by scale, etc.) is safe to say that is expected to be slow.

Hi Giovanni,

some clarifications: it's not just slow, we talk about hours (when it does not crash before). And I do not think it's normal, QGIS is able to calculate the centroids of all those geometries in less than 10 seconds if it does not perform the geometry validity check and in my opinion it makes no sense to make a validity check for labeling.

#3 - 2018-10-30 05:14 PM - Giovanni Manghi

*Hi Giovanni,
some clarifications: it's not just slow, we talk about hours (when it does not crash before). And I do not think it's normal, QGIS is able to calculate the centroids of all those geometries in less than 10 seconds if it does not perform the geometry validity check and in my opinion it makes no sense to make a validity check for labeling.*

I didn't explained myself very well: what i wanted to say is that I think (I may be wrong) that when the maps has tons of labels then qgis has always been slow (slow is the operation to draw labels, not to compute their position, I think). Have you tried on 2.18?

#4 - 2018-10-30 07:09 PM - Mario Baranzini

I didn't explained myself very well: what i wanted to say is that I think (I may be wrong) that when the maps has tons of labels then qgis has always been slow (slow is the operation to draw labels, not to compute their position, I think). Have you tried on 2.18?

Ok, but this is not the case. In fact, the labels to show are only 256 like the geometries. The 256 geometries are multipart and the parts are a lot (~100,000).

By placing the labels on the perimeter and not around the centroid, everything works perfectly. That's why I say that the problem is due to the calculation of centroids to place the labels and not the rendering of the labels.

On 2.18 there is exactly the same problem.

#5 - 2018-10-31 10:40 AM - Giovanni Manghi

On 2.18 there is exactly the same problem.

this prove my point, there are no geometry checks in 2.18 (the issue is not the check of geometries, is the drawing of a lot of labels that is historically slow in QGIS).

#6 - 2018-10-31 05:58 PM - Mario Baranzini

this prove my point, there are no geometry checks in 2.18 (the issue is not the check of geometries, is the drawing of a lot of labels that is historically slow in QGIS).

You're right on qgis 2.18, and in fact, by testing better, things are a bit different than what I said. Displaying the shapefile labels is slow, but not as much as in series 3 (a few minutes instead of hours), and the centroid algorithm works quickly (as in qgis 3 when the check is disabled).

But the problem for sure is not the drawing of the labels, in fact as written in a previous message, there are only 256 labels and if I create a layer with simple geometries (a point where are the centroids of the 256 geometries), with exactly the same labels of the shapefile at the same location, these labels are shown very quickly.

#7 - 2018-10-31 06:01 PM - Giovanni Manghi

Mario Baranzini wrote:

this prove my point, there are no geometry checks in 2.18 (the issue is not the check of geometries, is the drawing of a lot of labels that is historically slow in QGIS).

You're right on qgis 2.18, and in fact, by testing better, things are a bit different than what I said. Displaying the shapefile labels is slow, but not as much as in series 3 (a few minutes instead of hours), and the centroid algorithm works quickly (as in qgis 3 when the check is disabled).

But the problem for sure is not the drawing of the labels, in fact as written in a previous message, there are only 256 labels and if I create a layer with simple geometries (a point where are the centroids of the 256 geometries), with exactly the same labels of the shapefile at the same location, these labels are shown very quickly.

ok, so I think we have enough elements to say that this a regression, agree?

#8 - 2018-10-31 06:32 PM - Mario Baranzini

ok, so I think we have enough elements to say that this a regression, agree?

Yes, for labeling this is a regression. To display the same labels under the same conditions, it goes from 1 minute and 10 seconds in QGIS 2 to dozens of minutes with 100% cpu in QGIS 3.

#9 - 2018-11-01 09:30 AM - Giovanni Manghi

- Status changed from Feedback to Open
- Subject changed from QGIS takes a very long time to show labels in layers with a lot of multipart geometries to labelling is much slower in 3.4 compared to 2.18
- Description updated
- Priority changed from Normal to High
- Regression? changed from No to Yes

#10 - 2018-11-08 12:10 PM - Jürgen Fischer

- *Description updated*

#11 - 2018-11-09 03:56 AM - Nyall Dawson

- *Status changed from Open to In Progress*

#12 - 2018-11-09 08:13 PM - Nyall Dawson

- *Status changed from In Progress to Closed*

- *% Done changed from 0 to 100*

Applied in changeset commit:qgis|ff5a8bc0aebf51cd34c82ed0feaac21cbd22e2ce.