

QGIS Application - Bug report #8340

Very slow Spatialite DB creation on qgis 2.0/master

2013-07-22 06:05 AM - Giovanni Manghi

Status:	Closed		
Priority:	High		
Assignee:			
Category:	Data Provider/Spatialite		
Affected QGIS version:	master	Regression?:	No
Operating System:	Ubuntu	Easy fix?:	No
Pull Request or Patch supplied:	No	Resolution:	
Crashes QGIS or corrupts data:	No	Copied to github as #:	17124
Description			
<p>On qgis master (nightly+ubuntugis) on Ubuntu, the creation of a new SL db is very slow, it takes minutes, so long that at the beginning I was thinking it was a freeze.</p> <p>Does not occur on Windows and also on other Linux distros, like Debian (qgis self compiled).</p>			
Related issues:			
Duplicated by QGIS Application - Bug report # 9448: New spatialite layer neve...		Closed	2014-01-29

Associated revisions

Revision e04b426f - 2013-10-13 10:42 PM - Jürgen Fischer

spatialite: run InitSpatialMetadata(1) on SL >= 4.1 (fixes #8340)

History

#1 - 2013-07-24 04:52 PM - Daniel Vaz

It seems ok to me. But I am using Ubuntu 13.04 and qgis compiled from source.

#2 - 2013-07-25 04:48 AM - Giovanni Manghi

Daniel Vaz wrote:

| It seems ok to me. But I am using Ubuntu 13.04 and qgis compiled from source.

it seems specific of some library version, I'm on Ubuntu 12.04

#3 - 2013-08-08 07:35 AM - Giovanni Manghi

- Priority changed from Normal to Severe/Regression
- Operating System deleted (Linux/Ubuntu)

Now I see this also on a (clean) Windows/osgeo4w installation on both master and 1.8, while in the standalone 1.8 the creation of the SL is fast as expected.

#4 - 2013-08-11 05:21 PM - Giovanni Manghi

- Subject changed from *Very slow Spatialite DB creation on qgis-master/ubuntu* to *Very slow Spatialite DB creation on qgis master*

#5 - 2013-08-14 02:49 AM - Andreas Neumann

For me the registration of a new SL db takes not minutes, but 2-5 seconds. But still it is suboptimal to have seconds of unresponsiveness. Tested on OSGeo4W nightly and on Ubuntu 13.04 (self-compiled). But these are both rather fast machines.

#6 - 2013-08-14 02:51 AM - Giovanni Manghi

Andreas Neumann wrote:

For me the registration of a new SL db takes not minutes, but 2-5 seconds. But still it is suboptimal to have seconds of unresponsiveness. Tested on OSGeo4W nightly and on Ubuntu 13.04 (self-compiled). But these are both rather fast machines.

Hi Andreas, I have initially seen this on my Ubuntu machine and concluded it was a local issue, but then I started seeing this on my (clean) Windows VM that I use for testing with osgeo4w... I will check again.

#7 - 2013-08-15 05:53 PM - Daniel Vaz

Maybe it's a local issue like you said. I can't reproduce it here.

Please if you can provide some steps to follow, I will try to reproduce it.

#8 - 2013-08-15 06:27 PM - Giovanni Manghi

Daniel Vaz wrote:

Maybe it's a local issue like you said. I can't reproduce it here.
Please if you can provide some steps to follow, I will try to reproduce it.

no fancy steps to follow. I don't understand what I can have that is not ok, especially on Windows...

#9 - 2013-08-19 02:59 AM - Giovanni Manghi

- Resolution set to *invalid*
- Status changed from *Open* to *Closed*

I tested on another Windows machine and is ok. Still slow on my Linux machine, but not slow as before. So I guess that this is likely an issue with my pc.

#10 - 2013-09-27 09:28 AM - Paolo Cavallini

- Status changed from *Closed* to *Reopened*
- Priority changed from *Severe/Regression* to *High*
- Resolution deleted (*invalid*)

Confirmed here, on several machines, both Windows and Ubuntu. Unclear why it is slower on some machines than others. In worst cases it takes more than 10 minutes.

#11 - 2013-09-28 01:03 AM - Giovanni Manghi

- Target version set to Future Release - High Priority

Paolo Cavallini wrote:

Confirmed here, on several machines, both Windows and Ubuntu. Unclear why it is slower on some machines than others. In worst cases it takes more than 10 minutes.

and the resulting db is useless

see #8638

this ticket should be a blocker.

#12 - 2013-10-12 12:42 PM - Giovanni Manghi

- Subject changed from Very slow Spatialite DB creation on qgis master to Very slow Spatialite DB creation on qgis 2.0/master

#13 - 2013-10-13 10:11 AM - Jukka Rahkonen

Hi,

Metadata table "spatialite_history" gathers log data from all statements. For me with Windows Vista 32-bit and QGIS 2.0.1 the log looks like this:

```
spatial_ref_sys    table successfully created  2013-10-13T06:53:15.840Z
... snip ...
geom_cols_ref_sys  view 'geom_cols_ref_sys' successfully created  2013-10-13T06:53:23.913Z
spatial_ref_sys    table successfully populated  2013-10-13T07:11:34.246Z
```

Thus it took only 8 seconds to do everything that is needed except populating the spatial_ref_sys table and then populating the table took more than 18 minutes.

#14 - 2013-10-13 10:54 AM - Jukka Rahkonen

I think I found it. I could repeat the slow "InitSpatialMetadata" with Spatialite-gui and learned that the right way to do it with Spatialite 4.1.1 is as

```
select initspatialmetadata(1);
```

This takes 3 seconds with my computer. Similar thread from Spatialite users forum

https://groups.google.com/forum/#!msg/spatialite-users/La8BUrVKX_g/lGJKxnQzp1sJ

Jukka Rahkonen

#15 - 2013-10-13 12:02 PM - aperi2007 -

I guess the problem is due to a commit/transactional problem.

Infact I guess the actual code is simply call many insert without declare a transaction before of call it.

If is not declared a transaction the sqlite will add a transaction to every single insert.
So it is obviously slow.

As example:

this is slow:

```
begin transaction
  insert
end transaction
begin transaction
  insert
end transaction
begin transaction
  insert
end transaction
...
begin transaction
  insert
end transaction
```

Instead this is more fast:

```
begin transaction
  insert
  insert
  insert
  insert
end transaction
```

I guess it should be tried.

#16 - 2013-10-13 12:10 PM - aperi2007 -

I don't know the code,
but as reported from the last documentation from spatialite 4.1.1:

<http://www.gaia-gis.it/gaia-sins/spatialite-sql-4.1.0.html#p16>

if the optional argument transaction is set to TRUE the whole operation will be handled as a single Transaction (faster):
the default setting is transaction=FALSE (slower, but safer).
if the optional argument mode is not specified then any possible ESPG SRID definition will be inserted into the spatial_ref_sys table.
if the mode arg 'WGS84' (alias 'WGS84_ONLY') is specified, then only WGS84-related EPSG SRIDs will be inserted
if the mode arg 'NONE' (alias 'EMPTY') is specified, no EPSG SRID will be inserted at all

So the default method is more slower.

To have all in a single and fast transaction, is necessary to set "transaction=1":

```
InitSpatialMetaData( 1 )
```

#17 - 2013-10-13 12:41 PM - Even Rouault

For reference, this is also confirmed in OGR. See <http://trac.osgeo.org/gdal/ticket/5270> for the fix that has been applied.

#18 - 2013-10-13 02:05 PM - Jürgen Fischer

- *Status changed from Reopened to Closed*

Fixed in changeset commit:"e04b426f00f86a154ff74ed6bda5727086596b0f".