

QGIS Application - Feature request #9592

Oracle should only look in metadata table for metadata

2014-02-18 04:36 AM - Jonathan Moules

Status:	Open	
Priority:	Normal	
Assignee:		
Category:	Data Provider/Oracle	
Pull Request or Patch supplied:	No	Resolution:
Easy fix?:	No	Copied to github as #: 18172
Description When the "Only look in meta data table" checkbox is selected, QGIS should only use the MDSYS.USER_SDO_GEOM_METADATA for all of the metadata. This way it can get tablename, column and SRID. An option to skip the "type" check until the layer is added would also be useful. Currently it takes our system about 2 minutes to do a scan of all tables (i.e. from clicking "connect" to getting a full list of tables) because QGIS isn't trusting the USER_SDO_GEOM_METADATA.		

History

#1 - 2014-02-18 04:50 AM - Jürgen Fischer

because QGIS isn't trusting the USER_SDO_GEOM_METADATA.

no, just because USER_SDO_GEOM_METADATA doesn't carry information about the geometry type.

#2 - 2014-02-18 05:08 AM - Jonathan Moules

Jürgen Fischer wrote:

because QGIS isn't trusting the USER_SDO_GEOM_METADATA.

no, just because USER_SDO_GEOM_METADATA doesn't carry information about the geometry type.

Sorry I wasn't clear, I know that.

However, does QGIS **need** to know the Geometry type just to list the tables? I'm assuming it doesn't and thus I don't see why it **has** to find out the geometry type at that time.

Instead it could delay that check until the layer is added, which is what I'm asking for here.

I'm just guessing and may be wrong, so please feel free to correct me.

#3 - 2014-02-18 09:52 PM - Jukka Rahkonen

I would prefer using MDSYS.ALL_SDO_GEOM_METADATA. QGIS user is not necessarily the same who created the table.

About the geometry type check, how QGIS is trying to check it? As far as I know that info can only be found from index metadata and only if spatial index is created with "layer_gtype" parameter which is not compulsory. Or is it compulsory for QGIS?

#4 - 2014-02-19 02:09 AM - Jonathan Moules

Jukka Rahkonen wrote:

I would prefer using MDSYS.ALL_SDO_GEOM_METADATA. QGIS user is not necessarily the same who created the table.

You're quite right. I keep forgetting about that.

About the geometry type check, how QGIS is trying to check it? As far as I know that info can only be found from index metadata and only if spatial index is created with "layer_gtype" parameter which is not compulsory. Or is it compulsory for QGIS?

The Help says:

The table geometry is determined from the first 100 non-null geometry rows in the table.

Which is fine for one table; the problem I'm seeking to bypass here is to not perform that check on >500 tables consecutively when the user presses "connect".

#5 - 2014-02-19 02:30 AM - Nathan Woodrow

I just recently fixed (patched) this issue with the MS SQL driver. I just used the TOP 1 because just checking the first 100 is just as useless.

This is the fix that I applied commit:517b85937e. Made loading of the layers heaps quicker.

The other option is to have a geometry_columns query/table, which we have in the MS SQL and PostGIS providers. Some people don't like that because you have to maintain it but having a table like that really speeds up layer listing.

#6 - 2014-02-19 02:38 AM - Jonathan Moules

The other option is to have a geometry_columns query/table, which we have in the MS SQL and PostGIS providers. Some people don't like that because you have to maintain it but having a table like that really speeds up layer listing.

Yep, that came up in the discussion on the list, however:

- a) We already have to maintain a MAPINFO_TABLES and ArcSDE and those are troublesome enough.
- b) I figure the MDSYS.ALL_SDO_GEOM_METADATA should be sufficient for everything except data type if it is properly maintained (ours is).

I'm not averse to such a table, but I don't think it should be mandatory. My preferred option is to get what can be got from MDSYS.ALL_SDO_GEOM_METADATA, and not bother with geometry type until the layer is added to QGIS. I appreciate that won't be optimal for everyone, but if it's under the checkbox it should work quite well.

#7 - 2014-02-19 02:41 AM - Nathan Woodrow

I have had this discussion with others as well. I guess I'm the kind of person that likes seeing the geometry type in the dialog before I add the layer.

One issue with the MS SQL dialog is that you can't add until QGIS has found the geometry type. If that is the same case here maybe that is the issue.

Still load the geometry types in the background but still allow a user to add the layer before the type is known, once they click add go and work out the type then.

#8 - 2014-02-19 02:45 AM - Jonathan Moules

A further thought - does MS SQL allow multiple geometry types in the same column? I.e. mixing polygon, points etc, like you can in a TAB file, but can't in a shapefile? I suspect that's why Jürgen made it sample 100 rows and not just the first, because I believe Oracle does allow that (We're not using it to my knowledge, but still).

| *I guess I'm the kind of person that likes seeing the geometry type in the dialog before I add the layer.*

Oh yes, it's a nice to have, but if it's a choice between that and taking 2 minutes just to list the layers/tables, I'll take the mystery of not knowing. :-)

#9 - 2014-02-19 02:54 AM - Nathan Woodrow

Yes it does allow more than one geometry type, however if anyone does that they are dead to me :)

Taking the first sample 100 rows doesn't really work though because it might be points for the first 100 then 101 is a polygon. So I changed the MS SQL driver to just take the TOP 1 because in the best case, not mixed geometries, it will list everything really fast and it was you should have anyway. In the worst case it does the full group by and gives you all the types in the table, and if you have mixed geometries then that is punishment for doing so :)

| *Oh yes, it's a nice to have, but if it's a choice between that and taking 2 minutes just to list the layers/tables, I'll take the mystery of not knowing. :-)*

Yeah see but I don't think the issue is the checking the issue is that you can't load the layer until the check is done. Being to still able to add the layer while QGIS is off checking would fix that issue. Best of both words.

#10 - 2014-02-19 03:01 AM - Jonathan Moules

| *Yeah see but I don't think the issue is the checking the issue is that you can't load the layer until the check is done. Being to still able to add the layer while QGIS is off checking would fix that issue. Best of both words.*

I'm fine with the check being performed when the user adds a given layer (be it against 1 row, or 100) - that makes sense and is necessary for QGIS and only takes a fraction of a second. It's just that performing the check at the time of listing all layers (when the user clicks "connect") takes a long time because it has to do that fraction-of-a-second check 500 times. I don't doubt reducing the value to 1 will speed things up, but I suspect it would still be slow given enough tables.

#11 - 2014-02-19 03:04 AM - Nathan Woodrow

It's all done in a background thread though so it shouldn't block the UI, so you can have your cake and eat it too. List all the tables and check the geometry types while we wait if the user is slow to find what they need.

#12 - 2014-02-19 03:17 AM - Jonathan Moules

I'm not sure we're talking about the same thing (or maybe I'm mis-understanding how the internals work). The way it seems to work currently (from a users perspective) is that QGIS goes through every single table in MDSYS.ALL_SDO_GEOM_METADATA and gets the metadata for them one after the other.

There's a gauge at the bottom of the screen which shows what % of tables are complete and what the current table being checked is. Once the metadata has been collected for a table it is added to the list of tables and QGIS goes on to the next table. This entire process takes about 2 minutes for us. Thus if the table I want is near the bottom of the list, I have to wait almost 2 minutes before QGIS is showing it and I can add it.

What it sounds like you're talking about is:

- List all tables found in MDSYS.ALL_SDO_GEOM_METADATA immediately.
- While the user takes their time selecting which table they want, populate the "geometry type" field for each table.

That sounds better but isn't how the UI works currently, at least not for Oracle on our build (5192e8d).

#13 - 2014-02-19 03:22 AM - Nathan Woodrow

Ok looking at the code it seems that it's done different to the MS SQL provider. From the look of it it's still done in a thread but all the loading and checking is done there. In the MS SQL provider it will first get all the tables from sys.tables, or whatever it is, list them, then fire off a thread to look for the geometry columns. It's currently still semi blocking but I am fixing that in the next release.

#14 - 2014-03-04 05:56 AM - Jonathan Moules

I've just installed the latest nightly and it's still taking about 2 mins to list the tables so I guess this hasn't been implemented? (Checking the top 1 row)?

Although that said, if possible not getting the geometry when listing the tables would probably be optimal; only doing it when a table is added as a layer should speed this up hopefully.

#15 - 2014-03-10 05:57 AM - Jonathan Moules

Looking at qgsoracleprovider.cpp - it seems that Oracle is actually checking the first 1000 rows assuming I'm reading this line right:

```
SELECT DISTINCT gtype FROM (SELECT t.%1.sdo_gtype AS gtype FROM %2 t WHERE rownum<1000) WHERE rownum<=2
```

So I think there are several possible improvements:

1. Lower that to 1 per Nathan's suggestion
2. Use a threaded implementation like MS SQL apparently does.
3. Don't check geometry on listing - only when the layer is added.

#16 - 2015-10-01 11:38 PM - Johan Nel

The ideal solution is that in the logon dialog for the oracle we provide a schemaname. This way you can look in the USER_SDO_GEOM_METADATA with the schema owner. This allows you to log in as another user to a schema with reduced rights. In addition if you have a database with multiple schemas, eg, dev, test and prod, it will be easier to manage the information.

We would be happy to help with improving the implemetation

#17 - 2017-05-01 12:47 AM - Giovanni Manghi

- Easy fix? set to No

#18 - 2017-09-08 01:35 PM - Jürgen Fischer

- Assignee deleted (Jürgen Fischer)