

QGIS Application - Bug report #13336

Allow raster calculator with huge raster inputs

2015-09-09 07:03 AM - Simona Z

Status:	Closed	
Priority:	High	
Assignee:		
Category:	Raster Calculator	
Affected QGIS version:	3.4.1	Regression?: No
Operating System:		Easy fix?: No
Pull Request or Patch supplied:	No	Resolution: fixed/implemented
Crashes QGIS or corrupts data:	Yes	Copied to github as #: 21389
Description		
<p>Hello,</p> <p>I am trying to do some raster operation with Qgis' Raster Calculator but I get a minidump error and Qgis crashes (like minidump written to C:\Users\PROPRI~1\AppData\Local\Temp\qgis-20150909-155310-3892-7148-d20c5b7.dmp) every time. I am working with two raster files I created, so I would like to know if it's a problem with my rasters or with QGis</p> <p>Metadata of raster file (it's a binary file with 255 and 0 values):</p> <p>Driver GDAL provider GTiff GeoTIFF Dataset Description [raster pathname] AREA_OR_POINT=Area Band 1 STATISTICS_MAXIMUM=255 STATISTICS_MEAN=34.841265875025 STATISTICS_MINIMUM=0 STATISTICS_STDDEV=87.582012938486 Dimensions X: 18821 Y: 15837 Bands: 1 Origin 1.71397e+06,5.13445e+06 Pixel Size 2.00004,-2.00001 No Data Value -99999 Data Type Float64 - Sixty four bit floating point Pyramid overviews Layer Spatial Reference System +proj=tmerc +lat_0=0 +lon_0=9 +k=0.9996 +x_0=1500000 +y_0=0 +ellps=intl +towgs84=-104.1,-49.1,-9.9,0.971,-2.917,0.714,-11.68 +units=m +no_defs Layer Extent (layer original source projection) 1713966.4223982943221927,5102780.3500192686915398 : 1751609.1579853519797325,5134454.5267527680844069 Band Band 1 Band No 1</p>		

No Stats
No stats collected yet

Related issues:

Related to QGIS Application - Bug report # 17523: Raster calculator cannot ha...Closed2017-11-22

Associated revisions

Revision a665b7d7 - 2016-01-27 12:48 AM - Nyall Dawson

Don't crash when using raster calculator with huge raster (refs #13336)

Now we abort and advise the user if insufficient memory is available to perform the calculation.

At some future stage it would be nice to perform the calculations in blocks to allow this scenario, but for now it's better not to crash.

History

#1 - 2015-09-09 03:05 PM - Nyall Dawson

- Status changed from Open to Feedback

Can you share your raster file and the expression you are using for the calculation?

#2 - 2015-10-30 10:02 AM - Nicolas Cadieux

Hi Nyall,

Same problem with 2.12.0 on Windows 7-64

Files are large .tiff(84GB each) but this has worked in the past with 2.6 or 2.8. In my case, both are in the same projection, and I am not changing the output projection. This is simple math. layer1 - layer 2 (creating a normalized LiDAR layer). Same result using virtual raster sets of the same data.

Hope this helps.

Nicolas

#3 - 2015-11-06 05:18 AM - Giovanni Manghi

- Category set to Raster Calculator

- Priority changed from Normal to High

Nicolas Cadieux wrote:

Hi Nyall,

Same problem with 2.12.0 on Windows 7-64

Files are large .tiff(84GB each) but this has worked in the past with 2.6 or 2.8. In my case, both are in the same projection, and I am not changing the output projection. This is simple math. layer1 - layer 2 (creating a normalized LiDAR layer). Same result using virtual raster sets of the same data.

so, to replicate the issue is needed to use such big tiffs or it happens with any size ones? Any shareable dataset?

#4 - 2015-12-18 11:16 AM - Nicolas Cadieux

Hi,

I just saw the comments. Sorry for not responding. I thought I would get an email. I could share a file but it's 80 GB. I bit big but I could send it by ftp server. Or I could wait to find a smaller file that crashes QGIS. I did a test with a 250MB file and it work ok. It crashes with the big stuff. Send me an email to nicolas.cadieux at archeotec.ca if you need files to resolve this issues.

Nicolas

#5 - 2015-12-20 10:44 AM - Giovanni Manghi

- Subject changed from minidump error when using Raster calculator to QGIS crashes when using Raster calculator with huge raster inputs
- Status changed from Feedback to Open

Nicolas Cadieux wrote:

| Send me an email to nicolas.cadieux at archeotec.ca if you need files to resolve this issues.

would be happy to try test this huge rasters: giovanni dot manghi at naturalgis dot pt

#6 - 2015-12-20 07:37 PM - Nyal Dawson

Can you also send me a link to them? nyal dot dawson at gmail dot com

#7 - 2016-01-19 10:11 PM - Nyal Dawson

Did you manage to send this data? I'd like to address this for 2.14

#8 - 2016-01-20 01:48 PM - Nicolas Cadieux

Working on that file. For now, I can send you the dmp.

#9 - 2016-01-20 01:50 PM - Nicolas Cadieux

- File qgis-20160120-163811-3596-4868-b403532.7z added

trying again on the dmp. If it does not work, I will send it by email.

#10 - 2016-01-25 09:23 AM - Nicolas Cadieux

- File simulateRaster.7z added
- Assignee set to Nyal Dawson

Hi,

I have a solution. I have uploaded 2 files. A python 2.7 script that will create a simulated raster having the same general properties as my files. That will create a ASC file that has 145,857,241KB. (Will take time...) Then, use the gdal_translate file to save the ASC file as a tif. That file will have 84,013,848KB. With that, you can make a copy of this file (ex rename file1 and file2) and then do some simple raster math. It should crash as expected. (you will need to change the input and output files in these scripts (but as you are programs, you probably know this.)

Contact me if you have problems.

Nicolas

#11 - 2016-01-26 03:49 PM - Nyal Dawson

- Assignee deleted (Nyal Dawson)
- Tracker changed from Bug report to Feature request
- Subject changed from QGIS crashes when using Raster calculator with huge raster inputs to Allow raster calculator with huge raster inputs

#12 - 2016-01-26 03:49 PM - Nyal Dawson

Crash is now avoided, so moving this to a feature request - at some future stage we should perform the calculation in "blocks" to avoid the memory allocation errors.

#13 - 2016-03-23 03:25 PM - Drew Ricketts

Should this issue be resolved in QGIS 2.14? I'm trying to do a very simple raster calculation on a .tif using a mac, and QGIS crashes without giving a warning about memory.

#14 - 2016-03-28 02:52 PM - Nyal Dawson

Drew - sounds like a totally different issue.

#15 - 2016-06-09 03:07 AM - Jürgen Fischer

- Crashes QGIS or corrupts data changed from Yes to No
- Tracker changed from Feature request to Bug report

#16 - 2016-06-09 03:07 AM - Jürgen Fischer

- Tracker changed from Bug report to Feature request

#17 - 2016-07-14 09:26 AM - Giovanni Manghi

- Crashes QGIS or corrupts data changed from No to Yes
- Affected QGIS version changed from 2.10.1 to 2.16.0
- Target version set to Version 3.0
- Tracker changed from Feature request to Bug report

It still crashes, see #15281 (tested on a clean 2.16 installation on a Linux box with 16GB or ram).

#18 - 2016-09-03 03:53 AM - Johannes Kroeger

- File backtrace.log added

I guess I ran into this issue as well with a simple subtraction between two ~40000x40000 rasters.

Attached is a backtrace from gdb.

QGIS eats about 7GB of RAM while trying to do this, there would be 12GB more available for it though so I am not sure if it fills up something.

QGIS version 2.16.1-Nødebo.

It works perfectly fine if I do the same calculation with the gdal_calc.py utility outside QGIS.

If needed I can share the rasters, they are about 700MB each (compressed).

#19 - 2017-03-07 04:59 AM - Giovanni Manghi

- Status changed from Open to Feedback

Please try the NEW QGIS raster calculator available in the Processing toolbox of QGIS master (QGIS3).

#20 - 2017-03-07 04:59 AM - Giovanni Manghi

| If needed I can share the rasters, they are about 700MB each (compressed).

I would like to give them a try, thanks!

#21 - 2017-05-01 01:06 AM - Giovanni Manghi

- Easy fix? set to No

- Regression? set to No

#22 - 2017-05-26 05:21 PM - Giovanni Manghi

- Status changed from Feedback to Open

- Description updated

#23 - 2017-11-24 09:52 AM - Alessandro Pasotti

- Related to Bug report #17523: Raster calculator cannot handle huge raster files added

#24 - 2018-08-02 03:16 PM - Pawel Stankiewicz

Hi

Please change the affected version for Master and Regression for Yes because I've had to install 2.8 to make calculations for which there is not enough memory in 3.2.1 - 16 GB RAM, Windows 10, grid 50,553 x 32,083 = 1.6 Gpx, 6 GB not compressed tiff.

In the spring 3.0 started showing messages about not enough memory after Windows with running QGIS crashed, I think it was Windows fault - it crashes from time to time. Restoring a disk image restored also the calculator, reinstalling QGIS was not enough. Yesterday Windows with idle QGIS in the background had crashed, before I wanted to start using the calculator after a few months break.

I use Aoemi Backupper Standard for weekly disk images and I can explore files in disk images. I could restore last week image if it would be really necessary - it's time-consuming.

Maybe QGIS is too cautious to block calculations? I prefer to make the job done with a risk of crashes to no job done. It would be better to get an info about risk, to save changes, and than to risk.

Thanks for this great, usually, program.

#25 - 2018-08-03 10:47 AM - Giovanni Manghi

Pawel Stankiewicz wrote:

Hi

Please change the affected version for Master and Regression for Yes because I've had to install 2.8 to make calculations for which there is not enough memory in 3.2.1 - 16 GB RAM, Windows 10, grid 50,553 x 32,083 = 1.6 Gpx, 6 GB not compressed tiff.

interesting, any chance you can allow us download the dataset (and tell us the formula used)?

#26 - 2018-08-03 10:05 PM - Pawel Stankiewicz

Sure: <https://drive.google.com/file/d/1WNW-dlJcr83td4nuZSw9hQYJraDaP2II/view?usp=sharing>

Even grid * 1 formula doesn't work, but after choosing the grid extent. After choosing another layer extent the message isn't displayed, the calculator calculates, the result file has ~the correct size in MB, but with no data.

#27 - 2018-08-06 07:09 PM - Giovanni Manghi

- Affected QGIS version changed from 2.16.0 to 3.3(master)

Pawel Stankiewicz wrote:

Sure: <https://drive.google.com/file/d/1WNW-dlJcr83td4nuZSw9hQYJraDaP2II/view?usp=sharing>

Even grid * 1 formula doesn't work, but after choosing the grid extent. After choosing another layer extent the message isn't displayed, the calculator calculates, the result file has ~the correct size in MB, but with no data.

Confirmed:

here on Linux on master it always works (no crashes) but the results (with your dataset) is always a raster with only 0.

Wrong result, leaving the ticket as "corrupts data".

#28 - 2018-08-07 01:31 PM - Pawel Stankiewicz

Giovanni, do you confirm there is a problem or my description of the problem, which is a different from what you got? Previously I didn't write clearly enough: calculator calculates after selecting another **smaller** layer extent **from a different area**. 2.8 also makes no data in such case.

#29 - 2018-11-08 01:31 PM - Giovanni Manghi

- Status changed from Open to Feedback

Please try on QGIS 3.4.1, if the issue is still valid change the affected version, thanks.

#30 - 2018-11-08 04:45 PM - Nicolas Cadieux

Hi,

I have looked at this with 3.4.1. Raster math ends with the following message "Raster calculator : Insufficient memory available for operation." but without crashing. Problem still exist but does not end with a QGIS crash. This really needs to be addressed.

Nicolas

#31 - 2018-11-08 04:48 PM - Nicolas Cadieux

I don't see how to change the affected version. Can someone do this please?

Cheers

Nicolas

#32 - 2018-11-08 04:48 PM - Giovanni Manghi

I have looked at this with 3.4.1. Raster math ends with the following message "Raster calculator : Insufficient memory available for operation." but without crashing. Problem still exist but does not end with a QGIS crash. This really needs to be addressed.

r.mapcalculator has been re-introduced in GRASS7, this mean we will have it back in QGIS/Processing and we will then have again access to the most powerful raster calc tool (both in handling large data as also for the available functions/operators/logic).

#33 - 2018-11-08 04:49 PM - Giovanni Manghi

- Affected QGIS version changed from 3.3(master) to 3.4.1

#34 - 2018-11-08 04:59 PM - Nicolas Cadieux

Hi,

That is very good news. I have been using gdal_calc but it's limited to rasters having the same size, pixel posting and CRS. Is this the case with r.mapcalculator? Do you know if GPU or multi-threading is used by the GRASS module?

It would be nice to modify the Message to the following: "Raster calculator : Insufficient memory available for operation. " + "Operation requires X amount of memory. Use r.mapcalculator or gdal_calc found in processing.". This would informe the user that there is a viable option.

Thanks for the great work

Nicolas

#35 - 2018-11-08 05:00 PM - Giovanni Manghi

Nicolas Cadieux wrote:

*Hi,
That is very good news. I have been using gdal_calc but it's limited to rasters having the same size, pixel posting and CRS. Is this the case with r.mapcalculator?*

no

Do you know if GPU or multi-threading is used by the GRASS module?

not sure, you should ask the GRASS devs.

#36 - 2018-11-21 08:11 PM - Pawel Stankiewicz

Giovanni Manghi wrote:

| Please try on QGIS 3.4.1, if the issue is still valid change the affected version, thanks.

In 3.4.1 still not enough memory

#37 - 2018-12-24 10:39 AM - Alexander Bruy

Please test with latest master. There were some work to fix this issue.

#38 - 2018-12-27 09:00 PM - Pawel Stankiewicz

In 3.4.3 it works correctly. Thank you.

#39 - 2018-12-27 10:26 PM - Nyall Dawson

- Resolution set to fixed/implemented
- Status changed from Feedback to Closed

#40 - 2019-01-15 12:04 AM - Nicolas Cadieux

Looking good on 3.4.3! Congratulations

Files			
qgis-20160120-163811-3596-4868-b403532.7z	2.46 MB	2016-01-20	Nicolas Cadieux
simulateRaster.7z	648 Bytes	2016-01-25	Nicolas Cadieux
backtrace.log	2.74 KB	2016-09-03	Johannes Kroeger