

QGIS Application - Bug report #3735

GDAL stats not used

2011-04-14 04:24 AM - Redmine Admin

<b>Status:</b>	Closed	
<b>Priority:</b>	Low	
<b>Assignee:</b>	Redmine Admin	
<b>Category:</b>	Rasters	
<b>Affected QGIS version:</b>		<b>Regression?:</b> No
<b>Operating System:</b>	All	<b>Easy fix?:</b> No
<b>Pull Request or Patch supplied:</b>		<b>Resolution:</b> fixed
<b>Crashes QGIS or corrupts data:</b>		<b>Copied to github as #:</b> 13794
<b>Description</b>		
<p>I'm not sure if this behavior is expected or I am missing something.</p> <p>I have a 1 band raster dataset (an elevation DEM) in TIF format. It has statistics computed from "gdalinfo -stats mydata.tif" However, when I open it in QGIS (both 1.6 and 1.7 trunk), and look at the Metadata tab in Layer properties, there are two sections that contain band stats information. The first section is labeled "Band 1" and lists the same stats as found using gdalinfo. The second section is labeled "Band" and says "No stats collected yet" for band 1. QGIS has to create the stats when I apply a color map or stretch. For large rasters, this takes a while, and must be done each time it's used in QGIS.</p> <p>Has anyone experienced this behavior? Is this expected? Is there a way to let gdal compute all the stats that QGIS needs? Thanks.</p> <p>reported: <a href="http://lists.osgeo.org/pipermail/qgis-user/2011-March/011550.html">http://lists.osgeo.org/pipermail/qgis-user/2011-March/011550.html</a></p>		

History

#1 - 2011-04-15 09:48 AM - Tim Sutton

- Resolution set to wontfix
- Status changed from Open to Closed

Hi

The stats collected by gdal are only partially useful:

- they do not include stddev
- the min/max are estimated
- there is no mean
- there are various other stats not included

Because of this we don't use the gdal layer statistics and rather collect them ourselves. The process of collecting stats can be made more efficient by caching the stats in our own xml file similar to the aux.xml file - which is something we are planning for a future release.

For now I am closing this ticket as we will not be using GDAL stats in the foreseeable future.

Regards

Tim

**#2 - 2011-04-15 10:48 AM - Even Rouault**

Tim, I'm a bit surprised by your comments : `GDALComputeRasterStatistics()` **does** compute exact min, max, mean and stddev

**#3 - 2011-04-15 10:52 AM - diodata -**

Thanks for the explanation. Would it be possible then to store the computed raster statistics in an aux, xml or similar file, or internal db? QGIS needs to compute these stats every time that raster is opened. For rasters than are moderate size (1 GB) or larger, this could take a while. It feels strange when it needs to compute the same stats over and over again. Just a thought...

**#4 - 2011-04-15 10:54 AM - diodata -**

Never mind. I just read your comment again and see that you're planning to cache those stats. Thanks.

**#5 - 2011-04-15 04:55 PM - Tim Sutton**

- Status changed from Closed to Feedback
- Resolution deleted (wontfix)

Replying to [comment:2 rouault]:

| Tim, I'm a bit surprised by your comments : `GDALComputeRasterStatistics()` **does** compute exact min, max, mean and stddev

Hi Renoult

Thanks for your comments. Has this always been the case or has it changed in recent versions - I seem to remember that this was the reason I originally implemented my own stats gatherer and would be more than happy to drop QGIS's implementation in favour of GDAL's. I will investigate further and do just that if at all possible.

Regards

Tim

**#6 - 2011-04-16 12:58 AM - Even Rouault**

To the best of my knowledge, the implementation of `GDALComputeRasterStatistics()` has not been modified significantly in the last few GDAL versions. Just a few fixes here or there to ignore nodata values when computing the statistics. If you want exact stats (at the expense of a slower computation of course since GDAL needs to examine all pixels), you have to set the `bApproxOK` parameter to `FALSE`.

**#7 - 2011-08-04 03:26 AM - Tim Sutton**

- Pull Request or Patch supplied set to No
- Status changed from Feedback to Closed
- Resolution set to fixed

Fixed in commit:1c946725c1c6fdee05da6378bc42a0096318a2bf. Renoult thanks for your insight. For more info on the changes I have made, see:  
<http://linfiniti.com/2011/08/improvements-to-raster-performance-in-qgis-master/>

Thanks,

Tim