

Exiv2 - Bug #1042

Exiv2 nulls file on CIFS share when modifying Exif.Photo.UserComment

15 Mar 2015 18:51 - Calvin Browne

| | | | |
|---|-------------|------------------------|--------------------|
| Status: | Closed | Start date: | 15 Mar 2015 |
| Priority: | Normal | Due date: | |
| Assignee: | Robin Mills | % Done: | 100% |
| Category: | testing | Estimated time: | 1.00 hour |
| Target version: | 0.26 | | |
| Description | | | |
| Hi Guys, After running into this via shotwell, I set up a test CIFS directory of around 4000 jpg files. exiv2 -V exiv2 0.23 001700 (64 bit build) on Ubuntu 14.04.2 LTS (client & server). doing a for i in `ls`; do exiv2 -v -M"set Exif.Photo.UserComment charset=Ascii New Exif comment3" \$i; done will randomly, but reliably leave me with an empty file. The exiv2 command hangs and needs to be killed, and the file restored from backup ;-) Let me know what information I can give to help debug..... | | | |
| Related issues: | | | |
| Related to Exiv2 - Bug #1043: pyexiv2 fails on cifs shares on an Ubuntu client | | Closed | 19 Mar 2015 |

Associated revisions

Revision 3627 - 21 Mar 2015 16:35 - Robin Mills

#1042 and #1043. Don't use a Memlo object for small temporary files.

Revision 3630 - 24 Mar 2015 00:27 - Robin Mills

#1043 and #1042. Thanks to Thomas for showing that r3627 reintroduced #812. Thanks to Thoralf for suggesting msync Memlo fix.

Revision 3631 - 24 Mar 2015 10:36 - Robin Mills

#1042 #1043 #812 Added test regression detectors

Revision 3635 - 24 Mar 2015 22:25 - Robin Mills

#1043 #1042 #812. Thank You to Thomas for this "polishing" patch. Thank you to everybody who has worked on this issue. Adding all the comments on the three issues together comes to about 60 items by at least 6 contributors. And it involves platform issues, networking, Linux and Windows APIs. One of the most complex issues to arise in Exiv2. Well done everybody. And we've dealt with this quickly. Only 9 days since Calvin first reported #1042.

I use the term "complex" to mean many threads of technology. "complex" != "complicated". "complicated" = "difficult to understand". We try to avoid "complicated".

History

#1 - 15 Mar 2015 21:09 - Robin Mills

- Category set to testing
- Status changed from New to Assigned
- Assignee set to Robin Mills
- Target version set to 0.25

Thanks for reporting this, Calvin. Destroying your file isn't good. I'll have a look at this in the next few days.

Two thoughts come to mind:

1. Have you tried this on v0.24, or (even better) on our trunk?
2. File is mounted with CIFS. There was an issue with Windows/Virus checker locking files. <http://dev.exiv2.org/issues/984> I see you've said you're using Ubuntu server. Is there any possibility that another process is locking any of those files? For example, is shotwell or DigiKam busy indexing/thumbnailing files while you're updating all this metadata?

Random faults are always hard to diagnose. However I will try to reproduce this with Ubututu <-> Ubututu

Robin

#2 - 16 Mar 2015 00:05 - Robin Mills

Good News: I've set this up on a Mac Mini between Ubuntu 14.04 (client Virtual Maching on Mac) and Windows7 (server Virtual Machine on Mac). 3822 jpgs (all small files from my web site) - 358mb total. So an average of about 100k/file. No errors, no hangs using your command-line.

Tomorrow, I'll test:

1. Ubututu (Virtual Machine on Mac) <-> Ubuntu (Real machine) on Laptop.
2. Larger files

#3 - 16 Mar 2015 05:23 - Calvin Browne

Robin Mills wrote:

Thanks for reporting this, Calvin. Destroying your file isn't good. I'll have a look at this in the next few days.

Two thoughts come to mind:

1. Have you tried this on v0.24, or (even better) on our trunk?
2. File is mounted with CIFS. There was an issue with Windows/Virus checker locking files. <http://dev.exiv2.org/issues/984> I see you've said you're using Ubuntu server. Is there any possibility that another process is locking any of those files? For example, is shotwell or DigiKam busy indexing/thumbnailing files while you're updating all this metadata?

So - initially it manifested itself via shotwell, and I suspected some kind of locking issue. Which is why I went to the above.

The only possible other thing interacting with the files would be plex, but that's periodic, read only and direct on the server, and doesn't fit the behaviour.

Just to be clear, it's Ubuntu 14.04.2 LTS on both the CIFS client & server.

I'll see about getting 0.24 going and test that.

Random faults are always hard to diagnose. However I will try to reproduce this with Ubututu <-> Ubututu

Ok - it happens every time (ie before the end of the 4000 files). Which one it picks is random. It's been the first one, it's been the nth (where n is <4000), but it's been every time.

Thanks for the effort so far.

Robin

#4 - 16 Mar 2015 05:43 - Calvin Browne

```
~/exiv2-0.24/bin/exiv2 -V
exiv2 0.24 001800 (64 bit build)
```

```
for i in `ls`; do ~/exiv2-0.24/bin/exiv2 -v -M"set Exif.Photo.UserComment charset=Ascii New Exif comment4" $i; done
File 1/1: 1393827856454.jpg
Set Exif.Photo.UserComment "charset=Ascii New Exif comment4" (Comment)
File 1/1: 1393837729884.jpg
Set Exif.Photo.UserComment "charset=Ascii New Exif comment4" (Comment)
```

(hanging:

```
ls | 1393837729884.jpg
-rw-r--r-- 1 calvin 500 0 Mar 16 07:39 1393837729884.jpg
)
```

hmmmm.... now where do i find the trunk..... lemme look.

#5 - 16 Mar 2015 05:51 - Calvin Browne

Calvin Browne wrote:
<SNIP>

)
hmmmm.... now where do i find the trunk..... lemme look.

oh dear - no ./configure file with.....

svn checkout svn://dev.exiv2.org/svn/trunk

#6 - 16 Mar 2015 05:59 - Calvin Browne

Calvin Browne wrote:

Calvin Browne wrote:
<SNIP>

)
hmmmm.... now where do i find the trunk..... lemme look.

oh dear - no ./configure file with.....

svn checkout svn://dev.exiv2.org/svn/trunk

of course - reading the README helps ;-)

#7 - 16 Mar 2015 06:11 - Calvin Browne

~/exiv2-trunk/trunk/bin/exiv2 -Vexiv2 0.24 001800 (64 bit build)
Copyright (C) 2004-2013 Andreas Huggel.

About 50 lines.....

File 1/1: 20150131_192840.jpg
Set Exif.Photo.UserComment "charset=Ascii New Exif comment5" (Comment)

(hanging
ls | 20150131_192840.jpg
-rw-rw-r-- 1 calvin 500 0 Mar 16 08:09 20150131_192840.jpg)

:(

#8 - 16 Mar 2015 07:44 - Robin Mills

Calvin
There is almost no useful information in your last 4 or 5 reports. To be clear, I said I would set up and test Ubuntu <-> Ubuntu today and also investigate with larger files.

Can I infer that you have built the trunk and are using that now? And ls of hanging/empty file does not tell me anything new. Is this the first file, or after 500 files. Is it always the same file.

Please help me to help you. I'm going to have my breakfast (it's 7:42am in England) and perform the tests I said I would try. I'll report when I've done that.

#9 - 16 Mar 2015 08:05 - Calvin Browne

Robin Mills wrote:

Calvin
There is almost no useful information in your last 4 or 5 reports. To be clear, I said I would set up and test Ubuntu <-> Ubuntu today and also investigate with larger files.

Can I infer that you have built the trunk and are using that now? And ls of hanging/empty file does not tell me anything new. Is this the first file, or after 500 files. Is it always the same file.

Please help me to help you. I'm going to have my breakfast (it's 7:42am in England) and perform the tests I said I would try. I'll report when I've done that.

Apologies for not being very articulate.

Yup - I first tried the 2.4 tarball, and then the svn trunk.

The file truncation can happen on the first file, or the n-th file. It seems random.

#10 - 16 Mar 2015 09:19 - Robin Mills

I've reproduced this using the following setup. Client is a Linux VM running on the MacMini. Server is a laptop. Both are connected to the router by ethernet cable.

```
client: Linux rmillsmm-kubuntu 3.13.0-37-generic #64-Ubuntu SMP Mon Sep 22 21:28:38 UTC 2014 x86_64 x86_64 x86_64 GNU/Linux
server: Linux rmills-ubuntu 3.13.0-30-generic #54-Ubuntu SMP Mon Jun 9 22:45:01 UTC 2014 x86_64 x86_64 x86_64 GNU/Linux
```

As before, I've copied some images from my website to a test directory. 8000 images for a total size of 600mb. Average about 100k/image.

As you've said, it hangs (for me on the third file) and ^C breaks - level the file with length zero. Not Good!

```
1043 rmills@rmillsmm-kubuntu:~/smb4k/RMILLS-UBUNTU/rmills-ubuntu/home/rmills/temp/exiv2test $ for i in `ls`; do exiv2 -v -M"set Exif.Photo.UserComment charset=Ascii New Exif comment3" $i; done
File 1/1: 0001_2_3.jpg
Set Exif.Photo.UserComment "charset=Ascii New Exif comment3" (Comment)
File 1/1: 0001.jpg
Set Exif.Photo.UserComment "charset=Ascii New Exif comment3" (Comment)
File 1/1: 0011_2_3.jpg
Set Exif.Photo.UserComment "charset=Ascii New Exif comment3" (Comment)
^C
1044 rmills@rmillsmm-kubuntu:~/smb4k/RMILLS-UBUNTU/rmills-ubuntu/home/rmills/temp/exiv2test $
```

Here are the restored files:

```
1063 rmills@rmillsmm-kubuntu:~/smb4k/RMILLS-UBUNTU/rmills-ubuntu/home/rmills/temp/exiv2test $ ls -alt 001*.jpg
0001*.jpg
-rwxrwxr-x 1 rmills rmills 68998 Mar 16 08:08 0001_2_3.jpg
-rwxrwxr-x 1 rmills rmills 75647 Mar 16 08:08 0001.jpg
-rwxrwxr-x 1 rmills rmills 77930 Mar 16 08:08 0011_2_3.jpg
-rwxrwxr-x 1 rmills rmills 82040 Mar 16 08:08 0011.jpg
-rwxrwxr-x 1 rmills rmills 28702 Mar 16 08:04 001.jpg
1064 rmills@rmillsmm-kubuntu:~/smb4k/RMILLS-UBUNTU/rmills-ubuntu/home/rmills/temp/exiv2test $
```

There's nothing oddly different about "the breaker" 0011_2_3.jpg. It's much the same size as the 2 files which worked.

Here's what I'm thinking:

1. Is this peculiar to Ubuntu <-> Ubuntu CIFS, or can it impact any/every network
2. Is it Memory Mapping?

Concerning the first, I will try Ubuntu <-> Windows (smb/cifs) and Ubuntu <-> Ubuntu (NFS) and Ubuntu <-> Mac (Parallels magic and smb/cifs).

The second point is that we use Memory Mapping to perform I/O when it is available. When MM is not available, we use the "C" standard I/O fopen,ftell,fseek etc. Perhaps MM or the "switch to fopen" mechanism is not working correctly over CIFS. I'll build exiv2 and to always ignore MM and see what happens.

Please be patient and give me time to investigate this. I'm an unpaid volunteer and have other things in my life. I thank you for bringing this to my attention and assure you that I will work on this and report my findings.

#11 - 16 Mar 2015 09:43 - Calvin Browne

Glad you were able to reproduce - shout if there's something you want me to try - and I completely understand how this works, so no time pressure from my side ;-)

#12 - 16 Mar 2015 15:58 - Robin Mills

- Target version changed from 0.25 to 0.26

Calvin

Thanks for understanding. I've tested the following:

1. Linux <-> Linux with Memory Mapping Disabled. Hang. As earlier.
2. Linux <-> Windows. Given up. Can't get any smbclient on Ubuntu to accept my password for the Windows Laptop (I've tried smb4u, konqueror, dolphin and command-line)

3. Mac <-> Linux works
4. Mac <-> Windows works
5. Native - no networking. Mac/Linux/Windows works

These things take time. It's not like this is a simple code error. It's going to take some digging to find this. My suspect is the SMB client connection in Linux and perhaps we're failing to test some status somewhere. So - the data isn't ready - however we have ignored the error and charge on over the cliff.

I'm going to push this one to v0.26. I am rather busy at the moment and I suspect this issue has been round for a while. What's new here is your discovery.

If you have cycles to spend on this, perhaps you could try more combinations of client and server. I'm more suspicious of the client.

#13 - 18 Mar 2015 05:23 - Calvin Browne

Robin Mills wrote:

Calvin

Thanks for understanding. I've tested the following:

1. Linux <-> Linux with Memory Mapping Disabled. Hang. As earlier.
2. Linux <-> Windows. Given up. Can't get any smbclient on Ubuntu to accept my password for the Windows Laptop (I've tried smb4u, konqueror, dolphin and command-line)

Ok - so I resurrected my old windows xp laptop, shared a directory, mounted it under my linux client (`sudo mount -o uid=1000,gid=500 //192.168.2.145/test /media/test/ -t cifs`), copied a couple of hundred pics across and ran the command, and it hung and nulled the 29th file, and again about the 100th one.

Is there a windows binary for exiv2 that I should try this the other way round?
I'm going to try nfs and see what happens there...

#14 - 18 Mar 2015 10:51 - Robin Mills

Thank you for making the effort to test and report your result. Much appreciated. Together we can find and fix this matter.

I've build v0.23, v0.24 (both from the tar-ball on exiv2.org) and trunk this morning as MSVC 32bit static builds. No DLLs needed. I've zipped and posted them to: <http://clanmills.com/files/exiv2-32bit.zip> The trunk is the most interesting build. If that fails, it's unlikely that v0.24 and v0.23 will be OK.

I've provided those builds, so there's no need to run the command `exiv2 -vV` (verbose version) which reveals much more build information than plain `exiv2 -V`.

I did some googling into Ubuntu 14.04/Samba client connection. Lots of unhappy users - however this could be a coincidence. We keep digging for a while yet!

Thanks for working on this.

#15 - 18 Mar 2015 17:43 - Calvin Browne

Robin Mills wrote:

<http://clanmills.com/files/exiv2-32bit.zip> The trunk is the most interesting build. If that fails, it's unlikely that v0.24 and v0.23 will be OK.

Redirect loop :(

--Calvin

#16 - 18 Mar 2015 18:05 - Robin Mills

- File `exiv2-32bits.zip` added

How odd. It has indeed be removed from the server. I've restored it. And I've attached a copy.

#17 - 18 Mar 2015 18:54 - Calvin Browne

Robin Mills wrote:

How odd. It has indeed be removed from the server. I've restored it. And I've attached a copy.

got it - will test in the morning...
will also finish the nfs test...

#18 - 19 Mar 2015 10:13 - thoralf schulze

hi,

Robin Mills wrote:

1. Is it Memory Mapping?

probably related: <http://dev.exiv2.org/issues/1043>

according to man mount.cifs, cache=none aka directio "precludes mmaping files on this mount [...] Note too that no matter what caching model is used, the client will always use the pagecache to handle mmap'ed files. Writes to mmap'ed files are only guaranteed to be flushed to the server when msync() is called, or on close()." that might be relevant :-)

with kind regards,

t.

#19 - 21 Mar 2015 16:40 - Robin Mills

- Status changed from Assigned to Resolved

I'm going to mark this "Resolved". Further discussion of this issue will continue on [#1043](#).

#20 - 21 Aug 2015 10:41 - Robin Mills

- Status changed from Resolved to Closed

- Estimated time set to 1.00 h

This issue should have been closed in v0.25 and not left open for v0.26 (my error).

#21 - 22 Aug 2015 16:15 - Robin Mills

- % Done changed from 0 to 100

Files

| | | | |
|------------------|---------|-------------|-------------|
| exiv2-32bits.zip | 1.98 MB | 18 Mar 2015 | Robin Mills |
|------------------|---------|-------------|-------------|