

# CephFS - Bug #48203

## qa: quota failure

11/11/2020 08:44 PM - Patrick Donnelly

<b>Status:</b> Resolved	<b>% Done:</b> 0%
<b>Priority:</b> High	
<b>Assignee:</b> Luis Henriques	
<b>Category:</b>	
<b>Target version:</b> v16.0.0	
<b>Source:</b> Q/A	<b>ceph-qa-suite:</b>
<b>Tags:</b>	<b>Component(FS):</b> Client, kceph, qa-suite
<b>Backport:</b>	<b>Labels (FS):</b>
<b>Regression:</b> No	<b>Pull request ID:</b> 38112
<b>Severity:</b> 3 - minor	<b>Crash signature (v1):</b>
<b>Reviewed:</b>	<b>Crash signature (v2):</b>
<b>Affected Versions:</b>	
<b>Description</b>	
<pre>2020-11-04T18:03:28.550 INFO:tasks.workunit.client.1.smithi096.stderr:+ mv files limit/ 2020-11-04T18:03:28.559 INFO:tasks.workunit.client.1.smithi096.stderr:+ return 1 2020-11-04T18:03:28.561 DEBUG:teuthology.orchestra.run:got remote process result: 1 ... 2020-11-04T18:03:29.472 ERROR:teuthology.run_tasks:Saw exception from tasks. Traceback (most recent call last):   File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/run_tasks.py", line 90 , in run_tasks     manager = run_one_task(taskname, ctx=ctx, config=config)   File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/run_tasks.py", line 69 , in run_one_task     return task(**kwargs)   File "/home/teuthworker/src/github.com_batrick_ceph_wip-pdonnell-testing-20201103.210407/qa/task s/workunit.py", line 147, in task     cleanup=cleanup)   File "/home/teuthworker/src/github.com_batrick_ceph_wip-pdonnell-testing-20201103.210407/qa/task s/workunit.py", line 297, in _spawn_on_all_clients     timeout=timeout)   File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/parallel.py", line 84, in __exit__     for result in self:   File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/parallel.py", line 98, in __next__     resurrect_traceback(result)   File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/parallel.py", line 30, in resurrect_traceback     raise exc.exc_info[1]   File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/parallel.py", line 23, in capture_traceback     return func(*args, **kwargs)   File "/home/teuthworker/src/github.com_batrick_ceph_wip-pdonnell-testing-20201103.210407/qa/task s/workunit.py", line 425, in _run_tests     label="workunit test {workunit}".format(workunit=workunit)   File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/orchestra/remote.py", line 215, in run     r = self._runner(client=self.ssh, name=self.shortname, **kwargs)   File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/orchestra/run.py", lin e 446, in run     r.wait()   File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/orchestra/run.py", lin</pre>	

```
e 160, in wait
    self._raise_for_status()
File "/home/teuthworker/src/git.ceph.com_git_teuthology_master/teuthology/orchestra/run.py", lin
e 182, in _raise_for_status
    node=self.hostname, label=self.label
teuthology.exceptions.CommandFailedError: Command failed (workunit test fs/quota/quota.sh) on smit
hi096 with status 1: 'mkdir -p -- /home/ubuntu/ceph-test/mnt.1/client.1/tmp && cd -- /home/ubuntu/c
ephtest/mnt.1/client.1/tmp && CEPH_CLI_TEST_DUP_COMMAND=1 CEPH_REF=d2e2c4f1d55b90d4d72fec898522c82
d26aa11c4 TESTDIR="/home/ubuntu/ceph-test" CEPH_ARGS="--cluster ceph" CEPH_ID="1" PATH=$PATH:/usr/s
bin CEPH_BASE=/home/ubuntu/ceph-test/clone.client.1 CEPH_ROOT=/home/ubuntu/ceph-test/clone.client.1
adjust-ulimits ceph-coverage /home/ubuntu/ceph-test/archive/coverage timeout 3h /home/ubuntu/ceph-te
st/clone.client.1/qa/workunits/fs/quota/quota.sh'
```

From:  
/ceph/teuthology-archive/pdonnell-2020-11-04\_17:39:34-fs-wip-pdonnell-testing-20201103.210407-distro-basic-smithi/5590496/teuthology.log

See also discussion here: <https://tracker.ceph.com/issues/36593#note-13>

## History

---

### #1 - 11/11/2020 08:44 PM - Patrick Donnelly

- Description updated

### #2 - 11/11/2020 08:46 PM - Patrick Donnelly

In response to: <https://tracker.ceph.com/issues/36593#note-14>

Yes, there is not an easy solution here. I guess we have to transmit the truncate as soon as it occurs.

### #3 - 11/12/2020 03:27 PM - Luis Henriques

After discussing this with Jeff on the mailing-list<sup>1</sup> we agreed that the best thing to do is to simply revert to returning -EXDEV when doing cross quotarealms renames. Any other solution would have a big performance impact. This effectively means reverting commit dffdcd71458e ("ceph: allow rename operation under different quota realms").

[1] <https://www.spinics.net/lists/ceph-devel/msg49850.html>

### #4 - 11/13/2020 08:31 PM - Patrick Donnelly

Luis Henriques wrote:

After discussing this with Jeff on the mailing-list<sup>1</sup> we agreed that the best thing to do is to simply revert to returning -EXDEV when doing cross quotarealms renames. Any other solution would have a big performance impact. This effectively means reverting commit dffdcd71458e ("ceph: allow rename operation under different quota realms").

[1] <https://www.spinics.net/lists/ceph-devel/msg49850.html>

Also need to revert the userspace change too. That was done for issue [#39715](#).

To avoid forgetting why this is not a good idea, let's also add some good comments about why we don't allow cross-quota renames.

**#5 - 11/16/2020 11:02 AM - Luis Henriques**

Patrick Donnelly wrote:

Luis Henriques wrote:

After discussing this with Jeff on the mailing-list<sup>1</sup> we agreed that the best thing to do is to simply revert to returning -EXDEV when doing cross quotarealms renames. Any other solution would have a big performance impact. This effectively means reverting commit dffdc71458e ("ceph: allow rename operation under different quota realms").

[1] <https://www.spinics.net/lists/ceph-devel/msg49850.html>

Also need to revert the userspace change too. That was done for issue [#39715](#).

I can submit the revert for the fuse client, but I wonder if that's really necessary (other than for consistency between the two clients, of course). AFAICS the fuse client does not have the same problem because the truncate operations are sync. Or did you saw a similar failure?

To avoid forgetting why this is not a good idea, let's also add some good comments about why we don't allow cross-quota renames.

The kernel client revert is currently queued on the testing branch, and I tried to describe there the problem with the commit.

**#6 - 11/16/2020 05:32 PM - Patrick Donnelly**

Luis Henriques wrote:

Patrick Donnelly wrote:

Luis Henriques wrote:

After discussing this with Jeff on the mailing-list<sup>1</sup> we agreed that the best thing to do is to simply revert to returning -EXDEV when doing cross quotarealms renames. Any other solution would have a big performance impact. This effectively means reverting commit dffdc71458e ("ceph: allow rename operation under different quota realms").

[1] <https://www.spinics.net/lists/ceph-devel/msg49850.html>

Also need to revert the userspace change too. That was done for issue [#39715](#).

I can submit the revert for the fuse client, but I wonder if that's really necessary (other than for consistency between the two clients, of course). AFAICS the fuse client does not have the same problem because the truncate operations are sync. Or did you saw a similar failure?

The issue is not confined to a single client's consistent view of the quotas. We could have another client truncating files during the rename.

To avoid forgetting why this is not a good idea, let's also add some good comments about why we don't allow cross-quota renames.

The kernel client revert is currently queued on the testing branch, and I tried to describe there the problem with the commit.

Thanks!

**#7 - 11/16/2020 06:01 PM - Luis Henriques**

- Pull request ID set to 38112

Created pull-request <https://github.com/ceph/ceph/pull/38112> that simply reverts the fuse-client commit b8954e5734b3 ("client: optimize rename operation under different quota root").

**#8 - 12/02/2020 08:09 PM - Patrick Donnelly**

- Status changed from New to Resolved

- Component(FS) Client added

**#9 - 01/07/2021 03:50 PM - Patrick Donnelly**

- Status changed from Resolved to Need More Info

Hey Luis, I think this is still broken; the revert didn't work:

[https://pulpito.ceph.com/teuthology-2021-01-03\\_03:15:02-fs-master-distro-basic-smithi/5751916/](https://pulpito.ceph.com/teuthology-2021-01-03_03:15:02-fs-master-distro-basic-smithi/5751916/)

Can you take a look?

**#10 - 01/07/2021 04:42 PM - Luis Henriques**

Patrick Donnelly wrote:

Hey Luis, I think this is still broken; the revert didn't work:

[https://pulpito.ceph.com/teuthology-2021-01-03\\_03:15:02-fs-master-distro-basic-smithi/5751916/](https://pulpito.ceph.com/teuthology-2021-01-03_03:15:02-fs-master-distro-basic-smithi/5751916/)

Can you take a look?

Hmm... It looks like a distro kernel (kernel-4.18.0-240.1.1.el8\_3.x86\_64) is being used. Maybe the fix hasn't been backported to it? Or am I not looking at it right?

**#11 - 01/07/2021 05:57 PM - Patrick Donnelly**

*- Status changed from Need More Info to Resolved*

Luis Henriques wrote:

Patrick Donnelly wrote:

Hey Luis, I think this is still broken; the revert didn't work:

[https://pulpito.ceph.com/teuthology-2021-01-03\\_03:15:02-fs-master-distro-basic-smithi/5751916/](https://pulpito.ceph.com/teuthology-2021-01-03_03:15:02-fs-master-distro-basic-smithi/5751916/)

Can you take a look?

Hmm... It looks like a distro kernel (kernel-4.18.0-240.1.1.el8\_3.x86\_64) is being used. Maybe the fix hasn't been backported to it? Or am I not looking at it right?

Ah, good catch. That's why. Thanks for looking!