

## fs - Bug #13782

### Snapshotted files not properly purged

11/12/2015 12:48 PM - John Spray

<b>Status:</b>	Resolved	<b>Start date:</b>	11/12/2015
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>		<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>		<b>Affected Versions:</b>	
<b>Source:</b>	other	<b>ceph-qa-suite:</b>	
<b>Tags:</b>		<b>Component(FS):</b>	MDS
<b>Backport:</b>		<b>Labels (FS):</b>	
<b>Regression:</b>	No	<b>Pull request ID:</b>	
<b>Severity:</b>	3 - minor	<b>Crash signature:</b>	
<b>Reviewed:</b>			
<b>Description</b>			
Bad behaviour shown up by test TestStrays.test_snapshot_remove on this branch: <a href="https://github.com/ceph/ceph-qa-suite/tree/wip-snap-strays">https://github.com/ceph/ceph-qa-suite/tree/wip-snap-strays</a>			
It seems like there is a premature purge happening when the head revision is unlinked, the stray dentry for the file (file_a in the test) goes away at that point, and subsequently when the snapshot is deleted we are not purging anything. If you let it pass by the assertion that the purge shouldn't happen after unlink and run to the end of the test, you'll also see that the number of objects in the data pool does not fall to zero.			
<b>Related issues:</b>			
Related to fs - Bug #13777: Ceph file system is not freeing space		<b>Resolved</b>	<b>11/12/2015</b>

#### Associated revisions

##### Revision 511435f5 - 11/16/2015 01:17 PM - Yan, Zheng

client: avoid creating orphan object in Client::check\_pool\_perm()

Fixes: #13782

Signed-off-by: Yan, Zheng <[zyan@redhat.com](mailto:zyan@redhat.com)>

#### History

##### #1 - 11/12/2015 12:52 PM - John Spray

NB the difference between this and [#13777](#) is that in the other bug, we are seeing strays getting stuck (never getting purged) whereas here we are seeing purge go ahead but objects not being properly cleaned up

##### #2 - 11/12/2015 12:52 PM - John Spray

- Related to Bug #13777: Ceph file system is not freeing space added

##### #3 - 11/13/2015 08:59 AM - Zheng Yan

John Spray wrote:

It seems like there is a premature purge happening when the head revision is unlinked, the stray dentry for the file (file\_a in the test) goes away at that point, and subsequently when the snapshot is deleted we are not purging anything. If you let it pass by the assertion that the purge shouldn't happen after unlink and run to the end of the test, you'll also see that the number of objects in the data pool does not fall to zero.

this behaviour is expect. When modifying a regular inode with nlink == 1, we create a separate snap inode. If we delete unlink the head inode, it get purged. (snapshotted directory inode does not get purged until all snapshots are removed)

I ran commands in your test manually, number of objects in data pool did go to zero.

**#4 - 11/16/2015 12:10 PM - John Spray**

I've pushed an updated test branch which removes the mid-test assertions to avoid confusion.

Not only are the snapshot parts not purged, but there is also a HEAD instance of the file\_a 0th object visible at the end of the test.

**#5 - 11/16/2015 01:03 PM - Zheng Yan**

the head object is created by Client::check\_pool\_perm().(we should skip the check for reading snapshot). the reason of old snapshots not get purged is the same as [#13777](#). (mds does not have capability to remove snapshot)

**#6 - 12/01/2015 02:54 PM - John Spray**

- Status changed from New to Resolved

<https://github.com/ceph/ceph/pull/6603>

**#7 - 07/13/2016 05:54 AM - Greg Farnum**

- Component(FS) MDS added