

## Ceph - Bug #1194

### kclient: NFS reexport does not survive ceph fs remount

06/16/2011 02:23 PM - Brian Chrisman

<b>Status:</b>	Resolved	<b>% Done:</b>	0%
<b>Priority:</b>	Normal	<b>Spent time:</b>	0.00 hour
<b>Assignee:</b>	Brian Chrisman		
<b>Category:</b>			
<b>Target version:</b>	v0.38		
<b>Source:</b>		<b>Reviewed:</b>	
<b>Tags:</b>		<b>Affected Versions:</b>	
<b>Backport:</b>		<b>ceph-qa-suite:</b>	
<b>Regression:</b>	No	<b>Pull request ID:</b>	
<b>Severity:</b>	3 - minor	<b>Crash signature:</b>	
<b>Description</b>			
NFS doesn't survive restarts/remounts Reproduce: create new ceph fs export via NFS (on NFS client) nfs mount (on NFS client) create a couple files/directories stop nfs server umount ceph filesystem mount ceph filesystem start nfs server (on NFS client) access (cat) file in nfs mount  1. Stale NFS file handle  If I perform a 'find' on the ceph filesystem before attempting to access from NFS client after restart, no stale filehandle. Kernel client is still depending upon something being in a cache.  ceph version is ~0.28 era. kernel is 2.6.39 pulled within the last week from ceph-client git repo			

## History

### #1 - 06/16/2011 04:51 PM - Brian Chrisman

I reproduced the problem several times before submitting this bug but can't reproduce it now.  
I'm going to leave this open until tomorrow when I can try again with a fuller filesystem (only possible difference, though I'm pretty sure I ran a mkcephfs in there during my previous recreation).

### #2 - 06/20/2011 03:08 PM - Brian Chrisman

- File STALE\_on\_remount.messages.log added

- File STALE\_on\_remount.mds.log added

I was able to reproduce this.  
ESTALE shows up in the messages log.  
I don't see much in the mds log.

At this point, the fs is pretty much borked until I go into the ceph direct mount and perform lookup operations.  
I turned mds and kernel logging on just before running the operation I new would produce the ESTALE.  
I'm not sure how much I'd have to do to go from scratch as this is after copying a source tree to the filesystem with scp, then startup up the nfs export, then mounting, then attempting to build. That encounters an ESTALE and I continue to get the ESTALE after I've restarted everything from ceph on up (without umount/mount of the NFS client).

In this case, the NFS client is holding filehandles while the ceph/NFS server restart.

### #3 - 07/05/2011 09:22 AM - Sage Weil

- Target version set to v0.32

### #4 - 07/05/2011 11:08 AM - Sage Weil

- translation missing: en.field\_position set to 723

### #5 - 07/15/2011 08:58 PM - Sage Weil

- Assignee set to Sage Weil

I don't have things set up to reproduce/test this easily, but it looks like this is the problem. Can you give it a go?

```
diff --git a/fs/ceph/mds_client.c b/fs/ceph/mds_client.c
index 79743d1..5e0e8d1 100644
--- a/fs/ceph/mds_client.c
+++ b/fs/ceph/mds_client.c
@@ -1571,11 +1571,11 @@ static int set_request_path_attr(struct inode *rinode, struct dentry *rdentry,
     r = build_dentry_path(rdentry, ppath, pathlen, ino, freepath);
     dout(" dentry %p %llx/%.*s\n", rdentry, *ino, *pathlen,
         *ppath);
-     } else if (rpath) {
+     } else if (rpath || rino) {
         *ino = rino;
         *ppath = rpath;
         *pathlen = strlen(rpath);
-     dout(" path %.*s\n", *pathlen, rpath);
+     dout(" path #%llx/%.*s\n", rino, *pathlen, rpath);
     }

     return r;
```

### #6 - 07/23/2011 08:35 PM - Sage Weil

- Target version changed from v0.32 to v0.33

- translation missing: en.field\_position deleted (755)

- translation missing: en.field\_position set to 4

### #7 - 07/25/2011 09:23 AM - Sage Weil

- Subject changed from NFS reexport does not survive ceph fs remount to kclient: NFS reexport does not survive ceph fs remount

- translation missing: en.field\_position deleted (7)

- translation missing: en.field\_position set to 7

**#8 - 08/08/2011 09:18 AM - Sage Weil**

- Target version changed from v0.33 to v0.34

**#9 - 08/08/2011 09:21 AM - Sage Weil**

- translation missing: en.field\_position deleted (28)

- translation missing: en.field\_position set to 23

**#10 - 08/15/2011 01:31 PM - Sage Weil**

- Status changed from New to 7

pushed this to for-linus branch.

**#11 - 08/20/2011 11:08 AM - Sage Weil**

- Target version changed from v0.34 to v0.35

**#12 - 09/01/2011 01:06 PM - Sage Weil**

- Assignee changed from Sage Weil to Brian Chrisman

**#13 - 09/06/2011 09:49 PM - Sage Weil**

- Target version changed from v0.35 to v0.36

**#14 - 09/07/2011 12:03 PM - Sage Weil**

- translation missing: en.field\_position deleted (106)

- translation missing: en.field\_position set to 1

- translation missing: en.field\_position changed from 1 to 905

**#15 - 09/25/2011 02:14 PM - Sage Weil**

- Target version changed from v0.36 to v0.37

**#16 - 10/09/2011 08:39 PM - Sage Weil**

- Target version changed from v0.37 to v0.38

**#17 - 10/25/2011 04:46 PM - Sage Weil**

- Status changed from 7 to Resolved

going to assume the above fixed it until we hear otherwise :)

**Files**

---

STALE_on_remount.messages.log	916 KB	06/20/2011	Brian Chrisman
STALE_on_remount.mds.log	51.1 KB	06/20/2011	Brian Chrisman