

Linux kernel client - Bug #2392

First read of symlink after ceph filesystem mounted gives error

05/09/2012 04:55 PM - Mark Kirkwood

| | | | |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------|------------------------------|-----------|
| Status: | Resolved | % Done: | 0% |
| Priority: | High | Spent time: | 0.00 hour |
| Assignee: | Sage Weil | | |
| Category: | fs/ceph | | |
| Target version: | v3.5 | | |
| Source: | Community (user) | Reviewed: | |
| Tags: | | Affected Versions: | |
| Backport: | | ceph-qa-suite: | |
| Regression: | No | Crash signature (v1): | |
| Severity: | 3 - minor | Crash signature (v2): | |
| Description On client machine (Ubuntu 12.04): <pre>\$ mount -t ceph -o name=admin,secret=xxxx 192.168.1.101:6789:/ /mnt/osd \$ mkdir /mnt/osd/dir \$ ln -s /etc/hosts /mnt/osd/dir/hosts \$ umount /mnt/osd \$ mount -t ceph -o name=admin,secret=xxxx 192.168.1.101:6789:/ /mnt/osd \$ cat /mnt/osd/dir/hosts cat: /mnt/osd/dir/hosts: Invalid argument \$ cat /mnt/osd/dir/hosts ...contents of /etc/hosts as expected</pre> There are 3 other Ubuntu 12.04 hosts using the 0.46 ceph package, configured with mon, mts and osd on each host (see ceph.conf). | | | |

History

#1 - 05/14/2012 02:26 PM - Greg Farnum

I notice looking at your conf file that you have 3 MDSes. Are they all active? (ie, did you increase max_mds to 3)
If they are, see if you can reproduce with a single MDS. If they aren't, diagnosing this will be a lot easier with client and MDS logs.

#2 - 05/14/2012 04:37 PM - Mark Kirkwood

- File *ceph-logs-dev1.tar.gz* added
- File *kern.log* added

Ah - good point, no I had not updated max_mds. I redid the setup with 1 mds and 1 osd. Same issue, logs attached.

#3 - 05/14/2012 04:47 PM - Greg Farnum

- Assignee set to Greg Farnum

Okay, no guarantees but I will try and check this out at least briefly in the next day or two. :)

#4 - 05/15/2012 09:51 AM - Greg Farnum

Mark, can you repeat these with debug logging turned up? It'll take a fair bit of disk space but there's not very much in these logs. :)

#5 - 05/15/2012 04:26 PM - Mark Kirkwood

- File *ceph-logs-dev1.tar.gz* added

Sorry - remembering to enable debugging that would have been more helpful! Logs with debugging turned on attached.

#6 - 05/16/2012 11:45 AM - Greg Farnum

- Project changed from *Ceph* to *Linux kernel client*
- Category deleted (26)
- Assignee deleted (Greg Farnum)

Okay, this looks to me like it has to be a problem with the kernel client. The MDS definitely knows it's a symlink at all times, but I'm not the right person to look through the kernel code for this right now.

#7 - 05/16/2012 01:04 PM - Sage Weil

- Category set to *fs/ceph*
- Status changed from *New* to *12*

The problem is the lookup open intents stuff. We try to do a lookup + open, but it ends up that the lookup result is a symlink that you can't open; in fact the vfs needs to come back and follow the link.

Miklos' atomic open stuff is also broken, altho the error is "too many levesl of symbolic links".

I think this is just a matter of returning "success but not open" to the vfs.. not sure the right way to do that. Can probably just look at what nfs does.

#8 - 06/11/2012 09:50 AM - Sage Weil

- Assignee set to *Sage Weil*
- Priority changed from *Normal* to *High*
- Target version set to *v3.5*

This is going to be easy to fix once the atomic_open stuff is merged. Real Soon Now.

#9 - 07/30/2012 06:45 PM - Sage Weil

- Status changed from *12* to *7*

#10 - 08/10/2012 08:24 PM - Sage Weil

- Status changed from *7* to *Resolved*

Files

| | | | |
|-----------------------|-----------|------------|---------------|
| ceph.conf | 4.3 KB | 05/09/2012 | Mark Kirkwood |
| ceph-logs-dev1.tar.gz | 2.49 KB | 05/14/2012 | Mark Kirkwood |
| kern.log | 920 Bytes | 05/14/2012 | Mark Kirkwood |
| ceph-logs-dev1.tar.gz | 989 KB | 05/15/2012 | Mark Kirkwood |