

CephFS - Bug #23332

kclient: with fstab entry is not coming up reboot

03/13/2018 01:44 PM - Shreekara Shastry

Status: Closed	% Done: 0%
Priority: High	
Assignee: Zheng Yan	
Category:	
Target version: v14.0.0	
Source: Q/A	ceph-qa-suite:
Tags:	Component(FS): kceph
Backport:	Labels (FS):
Regression: No	Pull request ID:
Severity: 3 - minor	Crash signature (v1):
Reviewed:	Crash signature (v2):
Affected Versions:	
Description	
With 4 clients(2 fuse and 2 kernel clients), was running automated script for making fstab entry before reboot of clients . After reboot of each clients, I could see that client was not coming up ,after timeout of 600 seconds.	
Related issues:	
Related to CephFS - Bug #24879: mds: create health warning if we detect metad...	Resolved 07/11/2018

History

#1 - 03/13/2018 06:47 PM - Patrick Donnelly

- Subject changed from Client with fstab entry is not coming up reboot to kclient: with fstab entry is not coming up reboot
- Category deleted (Testing)
- Status changed from New to Need More Info
- Component(FS) kceph added
- Component(FS) deleted (ceph-fuse)

You probably have the wrong mons in the mount configuration.

#2 - 03/14/2018 03:44 AM - Zheng Yan

need add _netdev mount option

#3 - 03/14/2018 05:49 AM - Shreekara Shastry

Patrick Donnelly wrote:

You probably have the wrong mons in the mount configuration.

I have only one mon.

#5 - 03/16/2018 09:50 AM - Shreekara Shastry

- File `ceph-mds.ceph-sshreeka-run379-node4-mds.log` added
- File `ceph-mds.ceph-sshreeka-run379-node6-mds.log` added
- File `cierror.txt` added
- File `messages_ceph-sshreeka-run379-node5-client` added
- File `messages_ceph-sshreeka-run379-node7-client` added
- File `messages_ceph-sshreeka-run379-node8-client` added
- File `messages_ceph-sshreeka-run379-node9-client` added

#7 - 03/16/2018 05:29 PM - Vasu Kulkarni

Can someone please check the logs, we are having this issue 50% of the time in sanity runs.

#8 - 04/06/2018 10:28 PM - Patrick Donnelly

- Status changed from *Need More Info* to *New*
- Assignee set to *Zheng Yan*

#9 - 04/08/2018 01:12 AM - Zheng Yan

In `messages_ceph-sshreeka-run379-node5-client`

```
Mar 14 11:02:55 ceph-sshreeka-run379-node5-client mount: mount error 1 = Operation not permitted
Mar 14 11:02:55 ceph-sshreeka-run379-node5-client kernel: libceph: auth method 'x' error -1
```

looks like `fstab` didn't include correct secret

#10 - 04/08/2018 02:54 PM - Shreekara Shastry

Zheng Yan wrote:

In `messages_ceph-sshreeka-run379-node5-client`
[...]

looks like `fstab` didn't include correct secret

command `mount -a` works with same entries in `fstab`.

#11 - 04/19/2018 09:52 AM - Luis Henriques

Actually, the first failure seems to be a bit before:

```
Mar 14 08:10:00 ceph-sshreeka-run379-node5-client systemd: Stopping User Slice of root.
Mar 14 08:10:42 ceph-sshreeka-run379-node5-client kernel: ceph: mds1 caps stale
Mar 14 08:11:02 ceph-sshreeka-run379-node5-client kernel: ceph: mds1 caps stale
Mar 14 08:14:43 ceph-sshreeka-run379-node5-client kernel: libceph: mds0 172.16.115.100:6800 socket closed (con
state OPEN)
Mar 14 08:14:46 ceph-sshreeka-run379-node5-client chronyd[499]: Source 45.76.244.202 replaced with 64.113.44.5
4
Mar 14 08:14:47 ceph-sshreeka-run379-node5-client kernel: libceph: mds0 172.16.115.100:6800 connection reset
Mar 14 08:14:47 ceph-sshreeka-run379-node5-client kernel: libceph: reset on mds0
Mar 14 08:14:47 ceph-sshreeka-run379-node5-client kernel: ceph: mds0 closed our session
Mar 14 08:14:47 ceph-sshreeka-run379-node5-client kernel: ceph: mds0 reconnect start
Mar 14 08:14:47 ceph-sshreeka-run379-node5-client kernel: ceph: mds0 reconnect denied
Mar 14 08:15:20 ceph-sshreeka-run379-node5-client systemd-logind: Removed session 20.
Mar 14 08:15:20 ceph-sshreeka-run379-node5-client systemd: Removed slice User Slice of cepuser.
Mar 14 08:15:20 ceph-sshreeka-run379-node5-client systemd: Stopping User Slice of cepuser.
Mar 14 08:23:05 ceph-sshreeka-run379-node5-client kernel: libceph: auth method 'x' error -22
Mar 14 08:23:15 ceph-sshreeka-run379-node5-client kernel: libceph: auth method 'x' error -22
Mar 14 08:23:25 ceph-sshreeka-run379-node5-client kernel: libceph: auth method 'x' error -22
...
```

The client seems to be trying to renew caps while connection is being closed.

It's possible that this has been fixed in recent kernels and the fix hasn't been backported to your distro. Maybe it's worth trying a recent kernel and try to reproduce it.

#12 - 04/25/2018 06:06 PM - Patrick Donnelly

- Priority changed from Normal to High

- Target version set to v13.0.0

#13 - 04/25/2018 06:09 PM - Shreekara Shastry

Luis Henriques wrote:

Actually, the first failure seems to be a bit before:

[...]

The client seems to be trying to renew caps while connection is being closed.

It's possible that this has been fixed in recent kernels and the fix hasn't been backported to your distro. Maybe it's worth trying a recent kernel and try to reproduce it.

I've checked with the latest kernel version, still result was same

#14 - 04/26/2018 12:49 AM - Zheng Yan

kexec in dmesgs looks suspicious. client mounted cephfs, then used kexec to load kernel image again. All issues happened after kexec

#15 - 05/09/2018 09:05 PM - Patrick Donnelly

- Status changed from New to Need More Info
- Target version changed from v13.0.0 to v14.0.0
- ceph-qa-suite deleted (fs)

Zheng Yan wrote:

kexec in dmesgs looks suspicious. client mounted cephfs, then used kexec to load kernel image again. All issues happened after kexec

What do you need from Shreevara to verify this is the issue?

#16 - 05/10/2018 09:44 AM - Zheng Yan

I still don't think this is kernel issue. please patch kernel with below change and try again.

```
diff --git a/net/ceph/ceph_common.c b/net/ceph/ceph_common.c
index 584fdbef2088..3ca2d0c65d20 100644
--- a/net/ceph/ceph_common.c
+++ b/net/ceph/ceph_common.c
@@ -432,6 +432,7 @@ ceph_parse_options(char *options, const char *dev_name,
         err = -ENOMEM;
         goto out;
     }
+    printk("name: %s\n", opt->name);
     break;
     case Opt_secret:
         ceph_crypto_key_destroy(opt->key);
@@ -442,6 +443,7 @@ ceph_parse_options(char *options, const char *dev_name,
         err = -ENOMEM;
         goto out;
     }
+    printk("secret: %s\n", argstr[0].from);
     err = ceph_crypto_key_unarmor(opt->key, argstr[0].from);
     if (err < 0)
         goto out;
```

#17 - 07/12/2018 02:15 AM - Patrick Donnelly

- Related to Bug #24879: mds: create health warning if we detect metadata (journal) writes are slow added

#18 - 07/12/2018 02:15 AM - Patrick Donnelly

- Status changed from Need More Info to Closed

Closing this as it's a consequence of [#24879](#).

Files

reboot.jpg	343 KB	03/13/2018	Shreekara Shastry
ceph-mds.ceph-sshreeka-run379-node6-mds.log	56.4 KB	03/16/2018	Shreekara Shastry
cierror.txt	1.33 KB	03/16/2018	Shreekara Shastry
ceph-mds.ceph-sshreeka-run379-node4-mds.log	891 KB	03/16/2018	Shreekara Shastry
messages_ceph-sshreeka-run379-node7-client	323 KB	03/16/2018	Shreekara Shastry
messages_ceph-sshreeka-run379-node5-client	782 KB	03/16/2018	Shreekara Shastry
messages_ceph-sshreeka-run379-node8-client	467 KB	03/16/2018	Shreekara Shastry
messages_ceph-sshreeka-run379-node9-client	369 KB	03/16/2018	Shreekara Shastry