

## CephFS - Bug #24054

### kceph: umount on evicted client blocks forever

05/08/2018 10:37 PM - Patrick Donnelly

<b>Status:</b>	Resolved	<b>% Done:</b>	0%
<b>Priority:</b>	Urgent		
<b>Assignee:</b>	Zheng Yan		
<b>Category:</b>			
<b>Target version:</b>	v14.0.0		
<b>Source:</b>	Q/A	<b>Affected Versions:</b>	
<b>Tags:</b>		<b>ceph-qa-suite:</b>	
<b>Backport:</b>	mimic,luminous	<b>Component(FS):</b>	kceph
<b>Regression:</b>	No	<b>Labels (FS):</b>	
<b>Severity:</b>	3 - minor	<b>Pull request ID:</b>	
<b>Reviewed:</b>		<b>Crash signature:</b>	

#### Description

Failed test:

[/ceph/teuthology-archive/pdonnell-2018-05-08\\_01:06:46-kcephfs-mimic-testing-basic-smithi/2494030/teuthology.log](/ceph/teuthology-archive/pdonnell-2018-05-08_01:06:46-kcephfs-mimic-testing-basic-smithi/2494030/teuthology.log)

I had to kick it with `umount -f` which failed but that apparently caused the other umount (by the test) to finish:

```
pdonnell@smithi027:~$ mount
sysfs on /sys type sysfs (rw,nosuid,nodev,noexec,relatime)
proc on /proc type proc (rw,nosuid,nodev,noexec,relatime)
udev on /dev type devtmpfs (rw,nosuid,relatime,size=16393392k,nr_inodes=4098348,mode=755)
devpts on /dev/pts type devpts (rw,nosuid,noexec,relatime,gid=5,mode=620,ptmxmode=000)
tmpfs on /run type tmpfs (rw,nosuid,noexec,relatime,size=3281772k,mode=755)
/dev/sda1 on / type ext4 (rw,relatime,errors=remount-ro)
securityfs on /sys/kernel/security type securityfs (rw,nosuid,nodev,noexec,relatime)
tmpfs on /dev/shm type tmpfs (rw,nosuid,nodev)
tmpfs on /run/lock type tmpfs (rw,nosuid,nodev,noexec,relatime,size=5120k)
tmpfs on /sys/fs/cgroup type tmpfs (ro,nosuid,nodev,noexec,mode=755)
cgroup on /sys/fs/cgroup/systemd type cgroup (rw,nosuid,nodev,noexec,relatime,xattr,release_agent=/lib/systemd/systemd-cgroups-agent,name=systemd)
pstore on /sys/fs/pstore type pstore (rw,nosuid,nodev,noexec,relatime)
cgroup on /sys/fs/cgroup/net_cls,net_prio type cgroup (rw,nosuid,nodev,noexec,relatime,net_cls,net_prio)
cgroup on /sys/fs/cgroup/cpu,cpuacct type cgroup (rw,nosuid,nodev,noexec,relatime,cpu,cpuacct)
cgroup on /sys/fs/cgroup/perf_event type cgroup (rw,nosuid,nodev,noexec,relatime,perf_event)
cgroup on /sys/fs/cgroup/blkio type cgroup (rw,nosuid,nodev,noexec,relatime,blkio)
cgroup on /sys/fs/cgroup/freezer type cgroup (rw,nosuid,nodev,noexec,relatime,freezer)
cgroup on /sys/fs/cgroup/devices type cgroup (rw,nosuid,nodev,noexec,relatime,devices)
cgroup on /sys/fs/cgroup/memory type cgroup (rw,nosuid,nodev,noexec,relatime,memory)
cgroup on /sys/fs/cgroup/cpuset type cgroup (rw,nosuid,nodev,noexec,relatime,cpuset)
cgroup on /sys/fs/cgroup/hugetlb type cgroup (rw,nosuid,nodev,noexec,relatime,hugetlb)
hugetlbfs on /dev/hugepages type hugetlbfs (rw,relatime,pagesize=2M)
debugfs on /sys/kernel/debug type debugfs (rw,relatime)
systemd-1 on /proc/sys/fs/binfmt_misc type autofs (rw,relatime,fd=31,pgrp=1,timeout=0,minproto=5,maxproto=5,direct,pipe_ino=16613)
tracefs on /sys/kernel/debug/tracing type tracefs (rw,relatime)
mqueue on /dev/mqueue type mqueue (rw,relatime)
sunrpc on /run/rpc_pipefs type rpc_pipefs (rw,relatime)
nfsd on /proc/fs/nfsd type nfsd (rw,relatime)
fusectl on /sys/fs/fuse/connections type fusectl (rw,relatime)
configfs on /sys/kernel/config type configfs (rw,relatime)
/dev/mapper/vg_nvme-lv_5 on /var/lib/ceph type xfs (rw,relatime,attr2,inode64,noquota)
tmpfs on /run/user/1000 type tmpfs (rw,nosuid,nodev,relatime,size=3281772k,mode=700,uid=1000,gid=1
```

```
186)
172.21.15.200:6789,172.21.15.8:6790,172.21.15.8:6789:/ on /home/ubuntu/ceph/test/mnt.0 type ceph (r
w,relatime,name=0,secret=<hidden>,acl)
binfmt_misc on /proc/sys/fs/binfmt_misc type binfmt_misc (rw,relatime)
tmpfs on /run/user/1159 type tmpfs (rw,nosuid,nodev,relatime,size=3281772k,mode=700,uid=1159,gid=1
159)
pdonnell@smithi027:~$ umount -f /home/ubuntu/ceph/test/mnt.0
umount: only root can use "--force" option
pdonnell@smithi027:~$ sudo !!
sudo umount -f /home/ubuntu/ceph/test/mnt.0
umount: /home/ubuntu/ceph/test/mnt.0: target is busy
(In some cases useful info about processes that
use the device is found by lsof(8) or fuser(1).)
pdonnell@smithi027:~$ sudo umount -f /home/ubuntu/ceph/test/mnt.0
umount: /home/ubuntu/ceph/test/mnt.0: mountpoint not found
```

Fix should be similar to [#10915](#).

#### Related issues:

Follows CephFS - Bug #10915: client: hangs on umount if it had an MDS session...	<b>Resolved</b>	<b>02/19/2015</b>
Copied to CephFS - Backport #24188: luminous: kceph: umount on evicted client...	<b>Resolved</b>	
Copied to CephFS - Backport #24254: mimic: kceph: umount on evicted client bl...	<b>Resolved</b>	

#### History

##### #1 - 05/08/2018 10:38 PM - Patrick Donnelly

- Due date set to 02/20/2015
- Start date changed from 05/08/2018 to 02/20/2015
- Follows Bug #10915: client: hangs on umount if it had an MDS session evicted added

##### #2 - 05/14/2018 01:57 PM - Patrick Donnelly

- Status changed from New to Fix Under Review
- Backport set to luminous

<https://github.com/ceph/ceph/pull/21941>

##### #3 - 05/19/2018 09:57 AM - Zheng Yan

- Status changed from Fix Under Review to Pending Backport

##### #4 - 05/19/2018 10:04 AM - Nathan Cutler

- Copied to Backport #24188: luminous: kceph: umount on evicted client blocks forever added

##### #5 - 05/22/2018 01:47 PM - Patrick Donnelly

- Backport changed from luminous to mimic,luminous

##### #6 - 05/22/2018 06:15 PM - Patrick Donnelly

- Target version changed from v13.2.0 to v14.0.0

##### #7 - 05/23/2018 06:31 PM - Nathan Cutler

- Copied to Backport #24254: mimic: kceph: umount on evicted client blocks forever added

##### #8 - 10/22/2018 08:44 AM - Zheng Yan

- Status changed from Pending Backport to Resolved