

## CephFS - Bug #24133

### mds: broadcast quota to relevant clients when quota is explicitly set

05/15/2018 08:13 AM - Zhi Zhang

<b>Status:</b>	Resolved	<b>% Done:</b>	0%
<b>Priority:</b>	High		
<b>Assignee:</b>	Zhi Zhang		
<b>Category:</b>	Correctness/Safety		
<b>Target version:</b>	v14.0.0		
<b>Source:</b>	Community (dev)	<b>Affected Versions:</b>	
<b>Tags:</b>		<b>ceph-qa-suite:</b>	
<b>Backport:</b>	mimic,luminous	<b>Component(FS):</b>	MDS
<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

We found client won't get quota updated for a long time under following case. We found this issue on Luminous, but it might be also existed on previous versions.

#### Steps to reproduce:

1. mount a ceph-fuse client-1 (for admin) on host-1 and set quota as follows.

```
[ceph@x1 ~]$ sudo setfattr -n ceph.quota.max_bytes -v 100000 /mnt/cephfs/test10/  
[ceph@x1 ~]$ sudo setfattr -n ceph.quota.max_bytes -v 200000 /mnt/cephfs/test10/test11/
```

2. unmount above ceph-fuse client-1

3. mount another ceph-fuse client-2 (for user) on host-2 and copy some files into sub-directory test11. Since test10's quota is less than test11's, quota will be exceeded first on test10.

```
[ceph@x2 ~]$ sudo cp config /mnt/cephfs/test10/test11/2  
cp: error writing '/mnt/cephfs/test10/test11/2': Disk quota exceeded  
cp: failed to extend '/mnt/cephfs/test10/test11/2': Disk quota exceeded
```

4. mount a ceph-fuse client-1 (for admin) on host-1 again and change test10's quota.

```
[ceph@x1 ~]$ sudo setfattr -n ceph.quota.max_bytes -v 200000 /mnt/cephfs/test10/
```

5. unmount above ceph-fuse client-1

6. continue copying files into sub-directory test11 on host-2. It will still return quota exceeded for a long time, or until explicitly call getfattr to get quota back.

```
[ceph@x2 ~]$ sudo cp config /mnt/cephfs/test10/test11/3  
cp: error writing '/mnt/cephfs/test10/test11/3': Disk quota exceeded  
cp: failed to extend '/mnt/cephfs/test10/test11/3': Disk quota exceeded
```

The reason is that MDS only broadcasts quota under two cases, but above case is not one of them. So I think we could broadcast quota to relevant clients proactively if quota is explicitly set by someone.

**Related issues:**

Copied to CephFS - Backport #24205: luminous: mds: broadcast quota to relevan...	<b>Resolved</b>
Copied to CephFS - Backport #24206: mimic: mds: broadcast quota to relevant c...	<b>Resolved</b>

**History**

**#1 - 05/15/2018 08:18 AM - Zhi Zhang**

<https://github.com/ceph/ceph/pull/21997>

**#2 - 05/15/2018 07:54 PM - Patrick Donnelly**

- Assignee set to Zhi Zhang
- Priority changed from Normal to High
- Target version set to v13.2.0

**#3 - 05/15/2018 07:55 PM - Patrick Donnelly**

- Status changed from New to Fix Under Review

**#4 - 05/18/2018 10:51 PM - Patrick Donnelly**

- Target version changed from v13.2.0 to v14.0.0
- Backport set to mimic,luminous

**#5 - 05/20/2018 11:56 PM - Patrick Donnelly**

- Status changed from Fix Under Review to Pending Backport

**#6 - 05/21/2018 08:48 AM - Nathan Cutler**

- Copied to Backport #24205: luminous: mds: broadcast quota to relevant clients when quota is explicitly set added

**#7 - 05/21/2018 08:48 AM - Nathan Cutler**

- Copied to Backport #24206: mimic: mds: broadcast quota to relevant clients when quota is explicitly set added

**#8 - 08/27/2019 08:51 AM - Nathan Cutler**

- Status changed from Pending Backport to Resolved

While running with --resolve-parent, the script "backport-create-issue" noticed that all backports of this issue are in status "Resolved".