

fs - Bug #21412

cephfs: too many cephfs snapshots chokes the system

09/15/2017 09:53 PM - Wyllys Ingersoll

Status:	Closed	Start date:	09/15/2017
Priority:	Urgent	Due date:	
Assignee:	Zheng Yan	% Done:	0%
Category:	Correctness/Safety	Estimated time:	0.00 hour
Target version:	v13.0.0	Affected Versions:	
Source:	Community (user)	ceph-qa-suite:	
Tags:	snaps	Component(FS):	MDS
Backport:		Labels (FS):	
Regression:	No	Pull request ID:	
Severity:	2 - major		
Reviewed:			

Description

We have a cluster with /cephfs/.snap directory with over 4800 entries. Trying to delete older snapshots (some are over 6 months old on a pretty active file system) causes the "rmdir" command to hang, as well as any future operations on the .snap directory (such as 'ls'). Also, it is causing the number of blocked requests to grow indefinitely.

Ceph 10.2.7
Ubuntu 16.04.2
Kernel: 4.9.10

History

#1 - 09/15/2017 10:07 PM - Greg Farnum

- Project changed from Ceph to fs
- Category changed from 129 to 89

Can you dump the ops in flight on both the MDS and the client issuing the snap rmdir when this happens? And the perfcouners on the MDS?

My blind guess about what's blocking this is actually not snapshot trimming, but if the queue for deleting inodes (or one of the directory fragments, as you're on Jewel) is at its max size.

#2 - 09/16/2017 12:16 PM - Wyllys Ingersoll

Greg Farnum wrote:

Can you dump the ops in flight on both the MDS and the client issuing the snap rmdir when this happens? And the perfcouners on the MDS?

My blind guess about what's blocking this is actually not snapshot trimming, but if the queue for deleting inodes (or one of the directory fragments, as you're on Jewel) is at its max size.

What command(s) should I use to capture that info?

#3 - 09/17/2017 11:56 AM - Zheng Yan

- Assignee changed from Jos Collin to Zheng Yan

#4 - 09/17/2017 12:04 PM - Zheng Yan

ceph-mds.mds01.log.gz does not include useful information. The log was generated when mds replays log. Maybe the hang was caused by mds crash. does restarting mds resolve the hang?

#5 - 09/25/2017 01:41 PM - Greg Farnum

ceph daemon mds.<name> dump_ops_in_flight
ceph daemon mds.<name> perf dump

#6 - 09/25/2017 01:46 PM - Wyllys Ingersoll

Thanks. Im hesitant to trigger the issue again, last time it threw my cluster into major chaos that took several days to recover. Once I get data off of it, I will trigger the issue again and capture the info that you need.

#7 - 10/09/2017 01:53 PM - Wyllys Ingersoll

- File *perf_dump.after.txt* added

Here is data collected from a recent attempt to delete a very old and very large snapshot:

The snapshot extended attributes looks like:

```
1. file: cephfs/.snap/snapshot.2017-02-24_22_17_01-1487992621
   ceph.dir.entries="3"
   ceph.dir.files="0"
   ceph.dir.rbytes="30500769204664"
   ceph.dir.rctime="1504695439.09966088000"
   ceph.dir.retries="7802785"
   ceph.dir.rfiles="7758691"
   ceph.dir.rsubdirs="44094"
   ceph.dir.subdirs="3"
```

```
ops in flight during the deletion looks like: {
"ops": [],
"num_ops": 0
}
```

The problem is that it takes almost 24 hours to delete a single snapshot and it puts the cluster into a warning state whenever it is happening.

Is there a quicker "backdoor" way to purge our snapshots without blowing up the cluster? We really want to clean it up and get it back to a more usable state. At the current rate, it will literally take almost 13 YEARS to clean up the snapshots. Our only other alternative at this point is to destroy the entire filesystem and re-create it and then restore all of the data that was on it (we already backed it up, which took over a week).

#8 - 10/09/2017 01:55 PM - Wyllys Ingersoll

- File *dentry_lru.txt* added

Here is a dump of the cephfs 'dentry_lru' table, in case it is interesting.

#9 - 10/09/2017 01:56 PM - Wyllys Ingersoll

Note, the bug says "10.2.7" but we have since upgraded to 10.2.9 and the same problem exists.

#10 - 10/10/2017 02:51 AM - Zheng Yan

what do you mean "it takes almost 24 hours to delete a single snapshot"? 'rmdir .snap/xxx' tooks 24 hours or pgs on trimsnap states for 24 hours?

#11 - 10/10/2017 12:51 PM - Wyllys Ingersoll

The trimsnap states. The rmdir actually completes quickly, but the resulting operations throw the entire cluster into massive recovery storm that can takes days to recover from.

#12 - 04/09/2018 08:37 PM - Patrick Donnelly

- Subject changed from *too many cephfs snapshots chokes the system* to *cephfs: too many cephfs snapshots chokes the system*
- Category changed from *89* to *Correctness/Safety*
- Priority changed from *Normal* to *Urgent*
- Target version changed from *v10.2.10* to *v13.0.0*
- Source set to *Community (user)*
- Tags set to *snaps*
- Release deleted (*jewel*)
- Affected Versions deleted (*v10.2.7*)
- Component(*FS*) *MDS* added

Zheng, is this issue resolved with the snapshot changes for Mimic?

#13 - 04/11/2018 12:37 AM - Zheng Yan

this is actually osd issue. I talk to josh at cephalocon. He said it has already been fixed

#14 - 04/11/2018 12:37 AM - Zheng Yan

- Status changed from *New* to *Closed*

Files

ceph-mds.mds01.log.gz	557 KB	09/15/2017	Wyllys Ingersoll
perf_dump.after.txt	5.73 KB	10/09/2017	Wyllys Ingersoll
dentry_lru.txt	347 KB	10/09/2017	Wyllys Ingersoll