

Ceph - Bug #37568

CephFS remove snapshot result in slow ops

12/07/2018 10:01 AM - Francois Legrand

Status:	Resolved	Start date:	12/07/2018
Priority:	High	Due date:	
Assignee:	Zheng Yan	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	v14.0.0	Spent time:	0.00 hour
Source:	Community (dev)	Reviewed:	
Tags:		Affected Versions:	v13.2.2
Backport:	mimic,luminous	ceph-qa-suite:	fs
Regression:	No	Pull request ID:	25481
Severity:	3 - minor		

Description

Hello,

I have a ceph mimic cluster with cephfs.

I create few snapshots (mkdir .snap/test etc...) in different directories. So far so good.

But when I delete the snapshots (rmdir .snap/test etc...) the cluster get in a warn state with :

```
1. ceph -w
cluster:
id: 2fbbf089-a846-4c09-90bc-1dd9bd7af30f
health: HEALTH_WARN
3 slow ops, oldest one blocked for 11415 sec, mon.lpnceph01 has slow ops
...
2018-12-06 16:54:56.356518 mon.lpnceph-mon01 [WRN] Health check update: 3 slow ops, oldest one blocked for 11410 sec,
mon.lpnceph01 has slow ops (SLOW_OPS)
2018-12-06 16:55:05.856294 mon.lpnceph-mon01 [WRN] Health check update: 3 slow ops, oldest one blocked for 11415 sec,
mon.lpnceph01 has slow ops (SLOW_OPS)
2018-12-06 16:55:10.856657 mon.lpnceph-mon01 [WRN] Health check update: 3 slow ops, oldest one blocked for 11425 sec,
mon.lpnceph01 has slow ops (SLOW_OPS)
```

It's obviously related to the removal of snapshots because :

```
1. ceph daemon mon.lpnceph01 ops {
  "ops": [
    {
      "description": "remove_snaps({28=[3,4]} v0)",
      "initiated_at": "2018-12-06 13:44:41.396039",
      "age": 14549.148016,
      "duration": 14549.148028,
      "type_data": {
        "events": [
          {
            "time": "2018-12-06 13:44:41.396039",
            "event": "initiated"
          },
          {
            "time": "2018-12-06 13:44:41.396039",
            "event": "header_read"
          },
          {
            "time": "2018-12-06 13:44:41.396042",
            "event": "throttled"
          },
          {
            "time": "2018-12-06 13:44:41.396089",
            "event": "all_read"
          },
          {
            "time": "2018-12-06 13:44:41.396186",
            "event": "dispatched"
          },
          {
            "time": "2018-12-06 13:44:41.396190",
```

```

"event": "mon:_ms_dispatch"
},
{
"time": "2018-12-06 13:44:41.396191",
"event": "mon:dispatch_op"
},
{
"time": "2018-12-06 13:44:41.396192",
"event": "psvc:dispatch"
},
{
"time": "2018-12-06 13:44:41.396205",
"event": "osdmap:preprocess_query"
},
{
"time": "2018-12-06 13:44:41.396214",
"event": "osdmap:preprocess_remove_snaps"
},
{
"time": "2018-12-06 13:44:41.396220",
"event": "forward_request_leader"
},
{
"time": "2018-12-06 13:44:41.396258",
"event": "forwarded"
}
},
"info": {
"seq": 250448,
"src_is_mon": false,
"source": "mds.0 xxx.xxx.xxx.xxx:6800/2790459226",
"forwarded_to_leader": true
}
},
...

```

I tried to add in ceph.conf the lines :

```

[osd]
osd snap trim sleep = 0.6

```

as suggested in <http://lists.ceph.com/pipermail/ceph-users-ceph.com/2018-November/031227.html> but it doesn't solve the problem.

I had to restart the service :
systemctl restart [ceph-mon@lpnceph01.service](#)

to get the cluster back to healthy status.

Related issues:

Duplicated by Ceph - Bug #37782: Snapshot removal hangs	Duplicate	01/03/2019
Copied to Ceph - Backport #37693: mimic: CephFS remove snapshot result in slo...	Resolved	
Copied to Ceph - Backport #37694: luminous: CephFS remove snapshot result in ...	Resolved	

History

#1 - 12/10/2018 02:36 PM - Patrick Donnelly

- Subject changed from *Cephfs remove snapshot result in slow ops* to *CephFS remove snapshot result in slow ops*
- Assignee set to *Zheng Yan*
- Priority changed from *Normal* to *High*
- Target version set to *v14.0.0*

#2 - 12/11/2018 03:45 AM - Zheng Yan

- Project changed from *fs* to *Ceph*
- Category deleted (*89*)
- Status changed from *New* to *Need Review*
- Backport set to *mimic,luminous*
- Pull request ID set to *37568*

#3 - 12/11/2018 03:45 AM - Zheng Yan

- Pull request ID changed from 37568 to 25481

#4 - 12/18/2018 04:14 AM - Patrick Donnelly

- Status changed from Need Review to Pending Backport

#5 - 12/18/2018 11:10 AM - Nathan Cutler

- Copied to Backport #37693: mimic: CephFS remove snapshot result in slow ops added

#6 - 12/18/2018 11:10 AM - Nathan Cutler

- Copied to Backport #37694: luminous: CephFS remove snapshot result in slow ops added

#7 - 01/07/2019 10:07 PM - Patrick Donnelly

- Duplicated by Bug #37782: Snapshot removal hangs added

#8 - 02/22/2019 03:17 PM - Nathan Cutler

- Status changed from Pending Backport to Resolved