

CephFS - Bug #39437

osd: PriorityCache.cc: 265: FAILED ceph_assert(mem_avail >= 0)

04/23/2019 05:26 PM - Patrick Donnelly

Status:	Resolved	% Done:	0%
Priority:	Normal		
Assignee:	Mark Nelson		
Category:			
Target version:	v15.0.0		
Source:	Q/A	Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:		Component(FS):	
Regression:	No	Labels (FS):	
Severity:	3 - minor	Pull request ID:	27763
Reviewed:		Crash signature:	

Description

```
2019-04-20T18:30:11.748 INFO:tasks.ceph.osd.4.smithi079.stderr:/build/ceph-15.0.0-427-g609b64d/src
/common/PriorityCache.cc: In function 'void PriorityCache::Manager::balance()' thread 7f2025d78700
time 2019-04-20 18:30:11.747518
2019-04-20T18:30:11.748 INFO:tasks.ceph.osd.4.smithi079.stderr:/build/ceph-15.0.0-427-g609b64d/src
/common/PriorityCache.cc: 265: FAILED ceph_assert(mem_avail >= 0)
2019-04-20T18:30:11.750 INFO:tasks.ceph.osd.4.smithi079.stderr: ceph version 15.0.0-427-g609b64d (
609b64dbba750855b8c0d27110f536669da3205f) octopus (dev)
2019-04-20T18:30:11.750 INFO:tasks.ceph.osd.4.smithi079.stderr: 1: (ceph::__ceph_assert_fail(char
const*, char const*, int, char const*)+0x152) [0x8504d0]
2019-04-20T18:30:11.750 INFO:tasks.ceph.osd.4.smithi079.stderr: 2: (ceph::__ceph_assertf_fail(char
const*, char const*, int, char const*, char const*, ...) +0) [0x8506ab]
2019-04-20T18:30:11.750 INFO:tasks.ceph.osd.4.smithi079.stderr: 3: (PriorityCache::Manager::balanc
e()+0x417) [0xec69c7]
2019-04-20T18:30:11.750 INFO:tasks.ceph.osd.4.smithi079.stderr: 4: (BlueStore::MempoolThread::entr
y()+0x4ca) [0xdf1daa]
2019-04-20T18:30:11.751 INFO:tasks.ceph.osd.4.smithi079.stderr: 5: ((+0x76ba) [0x7f2033c196ba]
2019-04-20T18:30:11.751 INFO:tasks.ceph.osd.4.smithi079.stderr: 6: (clone()+0x6d) [0x7f203322041d]
```

/ceph/teuthology-archive/pdonnell-2019-04-17_06:07:08-multimds-wip-pdonnell-testing-20190417.032809-distro-basic-smithi/3857201/teuthology.log

Perhaps this can be resolved through tuning the OSD cache.

History

#1 - 04/24/2019 03:14 PM - Mark Nelson

Patrick Donnelly wrote:

[...]

/ceph/teuthology-archive/pdonnell-2019-04-17_06:07:08-multimds-wip-pdonnell-testing-20190417.032809-distro-basic-smithi/3857201/teuthology.log

Perhaps this can be resolved through tuning the OSD cache.

That assert is over-aggressive and ends up being thrown when there's not enough memory available to even add a single "chunk" of memory per cache for growth. Instead I'll just allow it and set mem_avail to zero. This will allow the OSD to run but may result in higher than osd_memory_target

memory usage in severely memory constrained scenarios.

Building a fix now.

Mark

#2 - 04/24/2019 03:42 PM - Patrick Donnelly

- Status changed from New to In Progress

- Assignee set to Mark Nelson

Mark Nelson wrote:

Patrick Donnelly wrote:

[...]

/ceph/teuthology-archive/pdonnell-2019-04-17_06:07:08-multimds-wip-pdonnell-testing-20190417.032809-distro-basic-smithi/3857201/teuthology.log

Perhaps this can be resolved through tuning the OSD cache.

That assert is over-aggressive and ends up being thrown when there's not enough memory available to even add a single "chunk" of memory per cache for growth. Instead I'll just allow it and set mem_avail to zero. This will allow the OSD to run but may result in higher than osd_memory_target memory usage in severely memory constrained scenarios.

Building a fix now.

Mark

Great! Thanks Mark.

#3 - 04/24/2019 06:34 PM - Mark Nelson

<https://github.com/ceph/ceph/pull/27763>

#4 - 04/24/2019 07:16 PM - Patrick Donnelly

- Status changed from *In Progress* to *Fix Under Review*

- Pull request ID set to 27763

#5 - 05/01/2019 08:26 PM - Patrick Donnelly

- Status changed from *Fix Under Review* to *Resolved*