

Ceph - Bug #9128

Newly-restarted OSD may suicide itself after hitting suicide time out value because it may need to search huge amount of objects

08/15/2014 03:53 AM - Zhi Zhang

Status:	Resolved	% Done:	0%
Priority:	Urgent	Spent time:	0.00 hour
Assignee:			
Category:			
Target version:			
Source:	Community (user)	Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:		Pull request ID:	
Regression:	No	Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			

Description

Stop one OSD daemon for a long time, like many hours even to 1 day, without marking it as out. During this time, there are still new writes to this cluster.

After restarting this OSD for a while, it may suicide itself because it often hits suicide time out value.(of course this depends on how long this OSD has been down and how many objects are written during this time)

If it has been down for enough long time and during this time there are lots of new writes, this issue is likely to happen. From the log and code, we see that OSD will search huge amount of objects in a for loop of PG::MissingLoc::add_source_info(...). At this time, the CPU of this OSD process is very high. So the health check will fail.

OSD log:

2014-08-04 12:58:30.486761 7f443802a700 10 osd.101 pg_epoch: 6115 pg[4.1acs0(v 5666'33417 lc 5537'30976 (5481'27975,5666'33417) local-les=6115 n=33192 ec=170 les/c 6103/6103 6114/6114/6114) [101,14,27,20,129,67,76,9,132,48,117] r=0 lpr=6114 pi=5550-6113/24 crt=5666'33417 lcod 0'0 mlcod 0'0 inactive m=2441 u=2441] search_for_missing b00bb1ac/default.5109.352_2bcb2a558999003fb691b35727c49984/head//4 5638'31694 is on osd.76(6)
2014-08-04 12:58:30.487781 7f4435825700 10 osd.101 pg_epoch: 6115 pg[4.249s0(v 5646'34066 lc 5537'31667 (5481'28667,5646'34066) local-les=6115 n=33843 ec=170 les/c 6103/6103 6114/6114/6114) [101,161,64,6,40,110,29,104,108,57,8] r=0 lpr=6114 pi=5534-6113/24 crt=5646'34066 lcod 0'0 mlcod 0'0 inactive m=2399 u=2399] search_for_missing fa0cc249/default.5106.441_b2d9436a4ff584e0a45978269e5a4dee/head//4 5638'33361 is on osd.104(7)

CPU usage:

top - 11:48:22 up 42 days, 1:38, 2 users, load average: 23.86, 31.97, 81.76
Tasks: 379 total, 1 running, 377 sleeping, 0 stopped, 1 zombie
Cpu(s): 49.3%us, 6.9%sy, 0.0%ni, 41.8%id, 0.0%wa, 0.0%hi, 2.0%si, 0.0%st
Mem: 48929020k total, 42767564k used, 6161456k free, 3016k buffers
Swap: 12582908k total, 0k used, 12582908k free, 160196k cached

PID	USER	PR	NI	VIRT	RES	SHR	S	%CPU	%MEM	TIME+	COMMAND
66399	yahoo		20	0 4133m	1.9g	208	S	317.9	4.0	813:58.52	ceph-osd
99827	yahoo		20	0 3698m	1.8g	92	S	200.8	3.9	397:58.54	ceph-osd
28381	yahoo		20	0 4121m	1.7g	0	S	101.2	3.6	414:07.98	ceph-osd
43089	yahoo		20	0 4069m	1.9g	716	S	100.6	4.1	405:23.48	ceph-osd

61566 yahoo 20 0 4038m 1.9g 852 S 100.6 4.0 445:19.62 ceph-osd

We can raise the op thread's time out value to further raise suicide time out value, but this just mitigates this issue.

Related issues:		
Related to Ceph - Bug #12523: osd suicide timeout during peering - search for...	Resolved	07/29/2015

Associated revisions

Revision 6aba0ab9 - 10/07/2014 12:49 AM - Wei Luo
Add reset_tp_timeout in long loop in add_source_info for suicide timeout
Fixes: #9128
Signed-off-by: luowei@yahoo-inc.com

History

#1 - 08/20/2014 09:19 AM - Sage Weil
- Status changed from New to 12
- Priority changed from Normal to High

sounds like we need to use the TPHandle and tp.reset_tp_handle() inside the search_For_missing loop

#2 - 09/09/2014 01:26 PM - Sage Weil
any progress on this?

#3 - 09/09/2014 06:35 PM - Guang Yang
Wei's patch - <https://github.com/ceph/ceph/pull/2371>

#4 - 09/10/2014 06:02 AM - Sage Weil
Guang Yang wrote:
Wei's patch - <https://github.com/ceph/ceph/pull/2371>

looks good, just needs a signed-off-by line

#5 - 10/07/2014 01:12 PM - Samuel Just
- Status changed from 12 to 7
- Priority changed from High to Urgent

#6 - 10/08/2014 11:04 AM - Samuel Just
- Status changed from 7 to Resolved