

Ceph - Bug #1594

pgs stuck degraded or active after 3 hours

10/03/2011 06:20 PM - Josh Durgin

Status:	Resolved	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:			
Category:	OSD		
Target version:	v0.38		
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature:	
Description			
From teuthology:~teuthworker/archive/nightly_coverage_2011-10-03/42/teuthology.log			
<pre>2011-10-03T18:04:41.494 INFO:teuthology.task.thrashosds.ceph_manager:2011-10-03 18:00:37.797592 pg v3327: 288 pgs: 1 active, 287 active+clean; 161 MB data, 170 GB used, 4002 GB / 4396 GB available; 2/114 degraded (1.754%) 2011-10-03 18:00:37.798535 mds e5: 1/1/1 up {0=0=up:active} 2011-10-03 18:00:37.798576 osd e106: 16 osds: 15 up, 15 in 2011-10-03 18:00:37.798674 log 2011-10-03 15:34:07.487919 mon.0 10.3.14.187:6791/0 55 : [INF] osd.3 out (down for 302.769562) 2011-10-03 18:00:37.798761 mon e1: 3 mons at {0=10.3.14.187:6791/0,1=10.3.14.168:6789/0,2=10.3.14.163:6790/0}</pre>			
osd and pg dumps are in teuthology:~ teuthworker/archive/nightly_coverage_2011-10-03/42/remote/ubuntu@sepia60.ceph.dreamhost.com/			
While this occurred, only one osd was down or out:			
<pre>osd.3 down out up_from 95 up_thru 95 down_at 97 last_clean_interval 37-93</pre>			
There's an active but not clean pg:			
<pre>0.19 2 2 2 0 8192 8388608 11232 11232 active 99'112 74'220 [8,10] [8,10] 0'0 2011-10-03 15:18:43.266093</pre>			

Associated revisions

Revision af6a9f30 - 10/07/2011 07:54 PM - Sage Weil

crush: try *all* bucket items when doing exhaustive search

N-1 isn't exhaustive.

Fixes: #1594

Signed-off-by: Sage Weil <sage@newdream.net>

History

#1 - 10/04/2011 04:39 PM - Josh Durgin

I reproduced this with debugging enabled. Logs are in vit:~joshd/thrash_stuck_active.
In this case there was 1 stuck active and 8 degraded with only one osd down and out.

#2 - 10/04/2011 09:40 PM - Sage Weil

Found one unrelated but, [a788bfdb93548751cec7184b65d42702cc207508](#).

I see one other possible badness:

- op is partially applied
 - osd.1 restarts, doesn't write it
 - osd.0 sees that it's missing
 - when the op is replayed, it recovers the object first before replying with dup
- ...but it isn't actually committed to disk on the target, only acked. we should probably be more strict here and don't ack pushes until it commits. it probably means being more aggressive about pipelining recovery operations, though, because the latency will go way up...

But anyway, the recovery is blocked because there are unfound objects on osd.0. Nothing really went *wrong* per se. :/

#3 - 10/07/2011 11:29 AM - Josh Durgin

Reproduced with 2.1p3 stuck in active since the up and acting sets were different. In this case 3 osds were marked out, with only 2 in.

Logs are in vit:/home/joshd/thrash_stuck_active2/. The osd dump was after restarting just the mons - all osds were up.

#4 - 10/09/2011 08:39 PM - Sage Weil

- Target version changed from v0.37 to v0.38

#5 - 10/09/2011 08:43 PM - Sage Weil

- translation missing: en.field_position set to 40

#6 - 10/12/2011 11:24 AM - Josh Durgin

- Status changed from New to Resolved

The bug in the second reproduced case was fixed by [af6a9f30696c900a2a8bd7ae24e8ed15fb4964bb](#).