

RADOS - Bug #22837

discover_all_missing() not always called during activating

01/31/2018 04:10 AM - David Zafman

Status:	Resolved	Start date:	01/30/2018
Priority:	Normal	Due date:	
Assignee:	David Zafman	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:		Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:	luminous	Component(RADOS):	
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature:	
Reviewed:			
Description			
<p>Sometimes discover_all_missing() isn't called so we don't get a complete picture of misplaced objects. This makes the new _update_calc_stats undercount misplaced as degraded which is the point of the changes. Also, this race makes a test case unreliable.</p> <p>On the run that does get the missing for osd.1 pg 1.0 goes from [1,0] -> [2,4] -> [2,4,3,5]</p> <p>A search_for_missing() was triggered during the [2,4] transition Then discover_all_missing during the [2,4,3,5] transition</p> <pre>2018-01-29 20:27:20.963 7f05606aa700 15 osd.2 pg_epoch: 35 pg[1.0(v 29'200 lc 0'0 (0'0,29'200] lo cal-lis/les=33/35 n=200 ec=26/26 lis/c 31/28 les/c/f 32/29/0 33/33/31) [2,4,3,5] r=0 lpr=33 pi=[28 ,33)/2 crt=29'200 mlcod 0'0 unknown m=200 u=200] build_might_have_unfound: built 0,1,3,4,5 2018-01-29 20:27:20.963 7f05606aa700 10 osd.2 pg_epoch: 35 pg[1.0(v 29'200 lc 0'0 (0'0,29'200] lo cal-lis/les=33/35 n=200 ec=26/26 lis/c 31/28 les/c/f 32/29/0 33/33/31) [2,4,3,5] r=0 lpr=33 pi=[28 ,33)/2 crt=29'200 mlcod 0'0 unknown m=200 u=200] discover_all_missing 200 missing, 200 unfound 2018-01-29 20:27:20.963 7f05606aa700 10 osd.2 pg_epoch: 35 pg[1.0(v 29'200 lc 0'0 (0'0,29'200] lo cal-lis/les=33/35 n=200 ec=26/26 lis/c 31/28 les/c/f 32/29/0 33/33/31) [2,4,3,5] r=0 lpr=33 pi=[28 ,33)/2 crt=29'200 mlcod 0'0 unknown m=200 u=200] discover_all_missing: osd.0: requesting pg_missin g_t 2018-01-29 20:27:20.963 7f05606aa700 10 osd.2 pg_epoch: 35 pg[1.0(v 29'200 lc 0'0 (0'0,29'200] lo cal-lis/les=33/35 n=200 ec=26/26 lis/c 31/28 les/c/f 32/29/0 33/33/31) [2,4,3,5] r=0 lpr=33 pi=[28 ,33)/2 crt=29'200 mlcod 0'0 unknown m=200 u=200] discover_all_missing: osd.1: requesting pg_missin g_t 2018-01-29 20:27:20.963 7f05606aa700 20 osd.2 pg_epoch: 35 pg[1.0(v 29'200 lc 0'0 (0'0,29'200] lo cal-lis/les=33/35 n=200 ec=26/26 lis/c 31/28 les/c/f 32/29/0 33/33/31) [2,4,3,5] r=0 lpr=33 pi=[28 ,33)/2 crt=29'200 mlcod 0'0 unknown m=200 u=200] discover_all_missing: osd.4: we already have pg_m issing_t</pre>			
Related issues:			
Related to Ceph - Bug #21803: objects degraded higher than 100%		Resolved	10/13/2017
Copied to RADOS - Backport #26992: luminous: discover_all_missing() not always...		Resolved	
History			
#1 - 02/01/2018 02:18 AM - David Zafman			
- Subject changed from discover_all_missing() not always called during peering to discover_all_missing() not always called during activating			
#2 - 02/03/2018 12:43 AM - David Zafman			
- Status changed from New to In Progress			

Part of <https://github.com/ceph/ceph/pull/20220>

#3 - 05/14/2018 09:38 PM - David Zafman

- Status changed from *In Progress* to *Resolved*

#4 - 08/20/2018 10:57 PM - David Zafman

- Related to Bug #21803: *objects degraded higher than 100% added*

#5 - 08/20/2018 11:19 PM - David Zafman

- Status changed from *Resolved* to *Pending Backport*

- Backport set to *luminous*

Based on information from <http://lists.ceph.com/pipermail/ceph-users-ceph.com/2017-October/021512.html> I'm marking this pending backport to luminous.

I can't say if this will be difficult to backport.

#6 - 08/21/2018 04:03 PM - Nathan Cutler

- Copied to Backport #26992: *luminous: discover_all_missing() not always called during activating added*

#7 - 09/25/2018 01:28 PM - Nathan Cutler

- Status changed from *Pending Backport* to *Resolved*