

RADOS - Bug #22881

scrub interaction with HEAD boundaries and snapmapper repair is broken

02/01/2018 06:09 PM - Sage Weil

Status:	Resolved	Start date:	02/01/2018
Priority:	Urgent	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			
symptom:			
<pre>2018-02-01T12:08:44.630 INFO:tasks.ceph.osd.4.smithi135.stderr:2018-02-01 12:08:44.620649 7f14491a 2700 -1 osd.4 pg_epoch: 752 pg[3.2(v 752'2572 (630'2174,752'2572] local-lis/les=751/752 n=50 ec=1 9/19 lis/c 751/751 les/c/f 752/752/0 726/751/572) [4,2,0] r=0 lpr=751 crt=752'2572 lcod 752'2570 m lcod 752'2570 active+clean+scrubbing] _scan_snaps no head for 3:6702c991::smithi09729659-280:194 (have MIN) 2018-02-01T12:08:44.632 INFO:tasks.ceph.osd.2.smithi097.stderr:2018-02-01 12:08:44.630462 7fd0eed2 1700 -1 osd.2 pg_epoch: 752 pg[3.2(v 752'2572 (653'2262,752'2572] local-lis/les=751/752 n=50 ec=1 9/19 lis/c 751/751 les/c/f 752/752/0 726/751/572) [4,2,0] r=1 lpr=751 luod=0'0 crt=752'2572 lcod 7 52'2570 active] _scan_snaps no head for 3:6702c991::smithi09729659-280:194 (have MIN) 2018-02-01T12:08:44.633 INFO:tasks.ceph.osd.0.smithi097.stderr:2018-02-01 12:08:44.631970 7fccfd0 5700 -1 osd.0 pg_epoch: 752 pg[3.2(v 752'2572 (630'2174,752'2572] local-lis/les=751/752 n=50 ec=1 9/19 lis/c 751/751 les/c/f 752/752/0 726/751/572) [4,2,0] r=2 lpr=751 luod=0'0 crt=752'2572 lcod 7 52'2570 active] _scan_snaps no head for 3:6702c991::smithi09729659-280:194 (have MIN)</pre>			
leading to			
<pre>Scrubbing terminated -- not all pgs were active and clean.</pre>			
scrub will accumulate a <code>cleaned_meta_map</code> that aligns to object boundaries that is used for the PG-type specific scrub step (<code>scrub_snapshot_metadata</code>). this is all good.			
however, for each chunk, which does not align to an object boundary (and should not!), we also do <code>_scan_snaps()</code> . This method does the snapmapper repair, and assumes that we have the head object in the same chunk in order to do it. this is not true in general and is basically broken.			
the problem unfortunately is that <code>_scan_snaps</code> happens locally on primary and replicas, while the <code>scrub_snapshot_metadata</code> (and accumulation of the <code>cleaned_meta_map</code>) happens only on primary. that means it cannot repair snapmapper on the replica.			
we probably need to accumulate <code>cleaned_meta_map</code> on replicas too and similarly gate the <code>_scan_snaps</code> on that...			
Related issues:			
Related to RADOS - Bug #23646: scrub interaction with HEAD boundaries and clo...		Resolved	04/10/2018
Related to RADOS - Bug #23909: snaps missing in mapper, should be: 188,18f,19...		Resolved	04/27/2018

Related to RADOS - Bug #24045: Eviction still raced with scrub due to preemption
Copied to RADOS - Backport #24016: luminous: scrub interaction with HEAD boun...

Resolved 05/07/2018
Resolved

History

#1 - 02/01/2018 06:09 PM - Sage Weil

- Backport set to luminous

#2 - 04/10/2018 09:06 PM - Sage Weil

- Subject changed from scrub interaction with HEAD boundaries is broken to scrub interaction with HEAD boundaries and snapmapper repair is broken

#3 - 04/10/2018 09:09 PM - Sage Weil

- Related to Bug #23646: scrub interaction with HEAD boundaries and clones is broken added

#4 - 04/20/2018 12:48 AM - David Zafman

- Status changed from Verified to In Progress
- Assignee set to David Zafman

#5 - 04/27/2018 03:02 PM - David Zafman

- Related to Bug #23909: snaps missing in mapper, should be: 188,18f,191,195,197,198,199,19d,19e,1a1,1a3 was r -2...repaired added

#6 - 05/04/2018 12:43 AM - David Zafman

- Status changed from In Progress to Pending Backport

#7 - 05/04/2018 12:46 AM - David Zafman

<https://github.com/ceph/ceph/pull/21546>

Backport the entire pull request which also fixes <http://tracker.ceph.com/issues/23909>

#8 - 05/04/2018 03:20 PM - Nathan Cutler

- Copied to Backport #24016: luminous: scrub interaction with HEAD boundaries and snapmapper repair is broken added

#9 - 05/08/2018 09:37 PM - David Zafman

- Related to Bug #24045: Eviction still raced with scrub due to preemption added

#10 - 05/25/2018 03:47 AM - David Zafman

- Status changed from Pending Backport to Resolved