

Ceph - Bug #8882

osd: osd tier remove ... leaves incomplete clones behind, confusing scrub

07/19/2014 09:30 AM - Sage Weil

Status:	Resolved	% Done:	0%
Priority:	Urgent	Spent time:	0.00 hour
Assignee:			
Category:	OSD		
Target version:			
Source:	Q/A	Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:		Pull request ID:	
Regression:		Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			
Description ubuntu@teuthology:/a/teuthology-2014-07-18_02:32:01-rados-master-testing-basic-plana/368480 rados/thrash/{clusters/fixed-2.yaml fs/btrfs.yaml msgr-failures/few.yaml thrashers/morepggrowse.yaml workloads/rados_api_tests.yaml}			
Related issues: Duplicated by Ceph - Bug #8881: scrub 85.0 cf2b2318/foo15/3/test-rados-api-pl... Duplicate 07/19/2014			

Associated revisions

Revision 956f2872 - 07/24/2014 05:07 PM - Sage Weil

osd/ReplicatedPG: do not complain about missing clones when INCOMPLETE_CLONES is set

When scrubbing, do not complain about missing clones when we are in a caching mode or when the INCOMPLETE_CLONES flag is set. Both are indicators that we may be missing clones and that that is okay.

Fixes: #8882

Signed-off-by: Sage Weil <sage@redhat.com>

Revision c42232e3 - 08/02/2014 12:00 AM - Sage Weil

osd/ReplicatedPG: do not complain about missing clones when INCOMPLETE_CLONES is set

When scrubbing, do not complain about missing clones when we are in a caching mode or when the INCOMPLETE_CLONES flag is set. Both are indicators that we may be missing clones and that that is okay.

Fixes: #8882

Signed-off-by: Sage Weil <sage@redhat.com>

(cherry picked from commit 956f28721dd98c5fb9eb410f4fe9e320b3f3eed3)

History

#1 - 07/23/2014 06:05 PM - Sage Weil

- Subject changed from scrub 85.5 2f0557ed/foo99/head/test-rados-api-plana76-2813-11/85 missing clones to osd: osd tier remove ... leaves incomplete clones behind, confusing scrub
- Status changed from New to 4

This is a stupid test with tiering teardown.

- set up cache pool, write a bunch of stuff
- some objects in teh cache tier are missing clones. this is normal, and scrub normally ignores it.
- test finishes, we osd tier remove cachepool, it clears the pool properties, and sets cache_mode = NONE
- scrub comes along and complains about missing clones because cache_mode == NONE
- the pool is deleted

A few ideas:

- 1- we leave cache_mode = WRITEBACK. i don't like this.
- 2- we set a flag on the pool that says 'incomplete clones' and set it when we set cache_mode to anything other than NONE. then change the scrub check to check for either mode != NONE or that flag (we have to behave with previously created cache pools)

I'm leaning toward #2?

#2 - 07/23/2014 06:24 PM - Sage Weil

- Status changed from 4 to Fix Under Review

#3 - 07/25/2014 10:32 AM - Sage Weil

- Status changed from Fix Under Review to Pending Backport

#4 - 08/12/2014 02:39 PM - Sage Weil

- Status changed from Pending Backport to Resolved

#5 - 06/29/2016 04:34 PM - Lazuardi Nasution

It seem that this issue still happen on Infernalis (v9.2.1). After removing the writeback cache by following procedure on <http://docs.ceph.com/docs/infernalis/rados/operations/cache-tiering/> the ex cachepool become has incomplete_clones flag and still leaves min_read_recency_for_promote and min_write_recency_for_promote flags. Beside, the ex storagepool become has lfor flag. Due to this condition (especially incomplete_clones flag on ex cachepool), I cannot use the same cachepool to build new cache tier with following error message.

Error ENOTEMPTY: tier pool {cachepool} has snapshot state; it cannot be added as a tier without breaking the pool

Anyway, I can delete this ex cachepool. Should I bring this issue as new bug report?

#6 - 06/29/2016 07:59 PM - Nathan Cutler

Should I bring this issue as new bug report?

The fix of this bug is already in infernalis, so yes, open a new bug.