

Ceph - Bug #23006

repair_test.yaml fails reproducibly in jewel integration testing

02/15/2018 11:16 AM - Nathan Cutler

Status:	Resolved	Start date:	02/15/2018
Priority:	Normal	Due date:	
Assignee:	David Zafman	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor		
Description			
Failure reason looks like this: "2018-02-07 09:08:23.888224 osd.3 172.21.15.148:6808/25279 13 : cluster [ERR] 5.0 : soid 5:a0216fbc:::repair_test_obj:head size 256 != size 1 from shard 3" in cluster log			
The failure is reproducible in the current 10.2.11 integration branch (wip-jewel-backports), and a similar failure appeared in a 10.2.6 integration run.			
See http://tracker.ceph.com/issues/21742#note-14 for the list of PRs included in the 10.2.11 integration branch.			
Logs 10.2.11:			
<ul style="list-style-type: none">http://pulpito.ceph.com/smithfarm-2018-02-06_21:07:15-rados-wip-jewel-backports-distro-basic-smithi/2160655/http://pulpito.ceph.com/smithfarm-2018-02-06_21:07:15-rados-wip-jewel-backports-distro-basic-smithi/2160751/http://pulpito.ceph.com/smithfarm-2018-02-15_10:48:02-rados-wip-jewel-backports-distro-basic-smithi/2191705/http://pulpito.ceph.com/smithfarm-2018-02-15_10:48:02-rados-wip-jewel-backports-distro-basic-smithi/2191704/			
The 10.2.6 failure text was ""2017-01-14 09:13:29.649276 osd.3 172.21.15.44:6800/865390 12 : cluster [ERR] 5.0 shard 3: soid 5:a0216fbc:::repair_test_obj:head size 1 != size 223 from auth oi 5:a0216fbc:::repair_test_obj:head(20'1 client.4352.0:1 dirty[data_digest omap_digest s 223 uv 1 dd 9a3a59aa od fffffff) in cluster log"			
Log 10.2.6:			
<ul style="list-style-type: none">http://pulpito.ceph.com/loic-2017-01-12_15:26:07-rados-wip-jewel-backports-distro-basic-smithi/712073/ (10.2.6 integration testing)			

History

#1 - 02/15/2018 11:17 AM - Nathan Cutler

@David, could you have a look? Are these two failures even related?

#2 - 02/15/2018 11:19 AM - Nathan Cutler

- Description updated

#3 - 02/15/2018 11:21 AM - Nathan Cutler

- Description updated

#4 - 02/15/2018 11:23 AM - Nathan Cutler

- Description updated

#5 - 02/15/2018 11:24 AM - Nathan Cutler

- Description updated

#6 - 02/15/2018 02:07 PM - Nathan Cutler

- Subject changed from *repair_test.yaml sometimes fails in jewel to repair_test.yaml fails reproducibly in jewel integration testing*
- Description updated

#7 - 02/16/2018 02:27 AM - David Zafman

When I do what I think is an equivalent truncation of a shard on current jewel, I get a single line showing all the errors:

```
2018-02-15 17:39:51.305625 7fef06ed5700 10 log_client logged 2018-02-15 17:39:49.264333 osd.1 127.0.0.1:6804/71530 14 : cluster [ERR] 1.0
shard 1: soid 1:602f83fe::foo:head data_digest 0x7c5a95e8 != data_digest 0xe850d09b from shard 0, data_digest 0x7c5a95e8 != data_digest
0xe850d09b from auth oi 1:602f83fe::foo:head(12'36 client.4114.0:36 dirty|data_digest|omap_digest s 147300440 uv 36 dd e850d09b od fffffff
alloc_hint [0 0]), size 1 != size 147300440 from auth oi 1:602f83fe::foo:head(12'36 client.4114.0:36 dirty|data_digest|omap_digest s 147300440 uv 36
dd e850d09b od fffffff alloc_hint [0 0]), size 1 != size 147300440 from shard 0
```

In the teuthology log of one of the runs in this tracker, we see 2 lines. The yaml only ignores lines with "size 1 != size"

```
2018-02-07T09:39:29.261 INFO:tasks.ceph.osd.3.smithi035.stderr:2018-02-07 09:39:29.263169 7f2086feb700 -1 log_channel(cluster) log [ERR] : 5.0
shard 3: soid 5:a0216fbc::repair_test_obj:head data_digest 0x6f0a661c != data_digest 0x3d1d6f0 from auth oi 5:a0216fbc::repair_test_obj:head(22'1
client.4257.0:1 dirty|data_digest|omap_digest s 258 uv 1 dd 3d1d6f0 od fffffff alloc_hint [0 0]), size 1 != size 258 from auth oi
5:a0216fbc::repair_test_obj:head(22'1 client.4257.0:1 dirty|data_digest|omap_digest s 258 uv 1 dd 3d1d6f0 od fffffff alloc_hint [0 0])
2018-02-07T09:39:29.261 INFO:tasks.ceph.osd.3.smithi035.stderr:2018-02-07 09:39:29.263178 7f2086feb700 -1 log_channel(cluster) log [ERR] : 5.0
: soid 5:a0216fbc::repair_test_obj:head data_digest 0x3d1d6f0 != data_digest 0x6f0a661c from shard 3, size 258 != size 1 from shard 3
```

I'm not very concerned about this, but need to look at the exact code some more.

#8 - 02/16/2018 04:52 PM - David Zafman

I do reproduce this with the exact sha1 (a95e2c4958b256d75ea1c732c2f2cce45f024081) you are testing. One fix is to backport 5f58301 however, that introduces another minor problem which is fixed in pull request [#20450](#).

I'm going to fix this by ignoring "size 258 != size" in yaml.

#9 - 02/16/2018 04:55 PM - David Zafman

- Status changed from *New to Need Review*
- Release set to *jewel*

<https://github.com/ceph/ceph/pull/20463>

#10 - 02/26/2018 08:08 PM - David Zafman

- Status changed from *Need Review to Resolved*