

Ceph - Bug #17194

Bluestore: Strange copy errors in osd log

09/01/2016 06:20 PM - Vitaliy Filippov

Status:	Won't Fix	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:			
Category:			
Target version:			
Source:	other	Affected Versions:	v10.2.2
Tags:	bluestore	ceph-qa-suite:	
Backport:		Pull request ID:	
Regression:	No	Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			

Description

Hi! I'm playing with a test setup of ceph jewel with bluestore and cephfs over erasure-coded pool with replicated pool as a cache tier. After writing some number of small files to cephfs I begin seeing the following error messages during the migration of data from cache to EC pool:

```
2016-09-01 10:19:27.364710 7f37c1a09700 -1 osd.0 pg_epoch: 329 pg[6.2cs0( v 329'388 (0'0,329'388] local-les=315 n=326 ec=279
les/c/f 315/315/0 314/314/314) [0,1,2] r=0 lpr=314 crt=329'387 lcod 329'387 mlcod 329'387 active+clean] process_copy_chunk data
digest 0x648fd38c != source 0x40203b61
```

```
2016-09-01 10:19:27.364742 7f37c1a09700 -1 log_channel(cluster) log [ERR] : 6.2cs0 copy from 8:372dc315:::200.0000002b:head
to 6:372dc315:::200.0000002b:head data digest 0x648fd38c != source 0x40203b61
```

These messages then repeat infinitely for the same set of objects with some interval. I'm not sure - does this mean some objects are corrupted in OSDs? (how to check?) Is it a bug at all?

History

#1 - 09/01/2016 09:16 PM - Samuel Just

- Subject changed from *Strange copy errors in osd log* to *Bluestore: Strange copy errors in osd log*

#2 - 09/01/2016 09:16 PM - Samuel Just

- Tags set to *bluestore*

#3 - 04/12/2017 10:52 PM - Sage Weil

- Status changed from *New* to *Won't Fix*

We haven't seen this in our testing on kraken or later; please reopen if you do!