

## Ceph - Bug #1526

### log bound mismatch after thrashing with bonnie

09/09/2011 09:57 AM - Josh Durgin

<b>Status:</b>	Resolved	<b>% Done:</b>	0%
<b>Priority:</b>	Normal	<b>Spent time:</b>	0.00 hour
<b>Assignee:</b>	Sage Weil		
<b>Category:</b>	OSD		
<b>Target version:</b>	v0.38		
<b>Source:</b>		<b>Reviewed:</b>	
<b>Tags:</b>		<b>Affected Versions:</b>	
<b>Backport:</b>		<b>ceph-qa-suite:</b>	
<b>Regression:</b>	No	<b>Pull request ID:</b>	
<b>Severity:</b>	3 - minor	<b>Crash signature:</b>	
<b>Description</b>			
A few thrashosds runs showed this yesterday, but this one had the most (more in teuthology:~teuthworker/archive/thrash_tests_2011-09-08_2/73/ceph.log):			
2011-09-08 18:38:01.166610 osd0 10.3.14.169:6800/1106 53 : [ERR] 1.0 log bound mismatch, info (223 '14,329'66] actual [3'1,329'66]			
2011-09-08 18:38:06.925916 osd0 10.3.14.169:6800/1106 54 : [ERR] 1.0 log bound mismatch, info (223 '14,360'67] actual [3'1,360'67]			
2011-09-08 18:40:14.893511 osd0 10.3.14.169:6800/1106 55 : [ERR] 1.b log bound mismatch, info (174 '4,419'34] actual [3'1,419'34]			
2011-09-08 18:40:28.316841 osd0 10.3.14.169:6800/1106 56 : [ERR] 1.b log bound mismatch, info (174 '4,419'34] actual [3'1,419'34]			
2011-09-08 18:40:33.998139 osd0 10.3.14.169:6800/1106 57 : [ERR] 1.b log bound mismatch, info (174 '4,419'34] actual [3'1,419'34]			
2011-09-08 18:40:39.663927 osd0 10.3.14.169:6800/1106 58 : [ERR] 1.b log bound mismatch, info (174 '4,419'34] actual [3'1,419'34]			
2011-09-08 18:41:03.087029 osd0 10.3.14.169:6800/1106 59 : [ERR] 1.a log bound mismatch, info (261 '40,467'79] actual [261'24,467'79]			
2011-09-08 18:41:03.087300 osd0 10.3.14.169:6800/1106 60 : [ERR] 1.b log bound mismatch, info (174 '4,419'34] actual [3'1,419'34]			
2011-09-08 18:41:14.130318 osd0 10.3.14.169:6800/1106 61 : [ERR] 1.a log bound mismatch, info (261 '40,467'79] actual [261'24,467'79]			
2011-09-08 18:41:14.130748 osd0 10.3.14.169:6800/1106 62 : [ERR] 1.b log bound mismatch, info (174 '4,419'34] actual [3'1,419'34]			
2011-09-08 18:41:25.102701 osd0 10.3.14.169:6800/1106 63 : [ERR] 1.a log bound mismatch, info (261 '40,467'79] actual [261'24,467'79]			
2011-09-08 18:41:25.102982 osd0 10.3.14.169:6800/1106 64 : [ERR] 1.b log bound mismatch, info (174 '4,419'34] actual [3'1,419'34]			
2011-09-08 18:41:49.027975 osd0 10.3.14.169:6800/1106 65 : [ERR] 1.a log bound mismatch, info (261 '40,467'79] actual [261'24,467'79]			
2011-09-08 18:41:49.028258 osd0 10.3.14.169:6800/1106 66 : [ERR] 1.b log bound mismatch, info (174 '4,419'34] actual [3'1,419'34]			
2011-09-08 18:41:56.752982 osd0 10.3.14.169:6800/1106 67 : [ERR] 1.a log bound mismatch, info (261 '40,467'79] actual [261'24,467'79]			
2011-09-08 18:41:56.753306 osd0 10.3.14.169:6800/1106 68 : [ERR] 1.b log bound mismatch, info (174 '4,419'34] actual [3'1,419'34]			
2011-09-08 18:42:01.557027 osd0 10.3.14.169:6800/1106 69 : [ERR] 1.a log bound mismatch, info (261 '40,467'79] actual [261'24,467'79]			
2011-09-08 18:42:01.557296 osd0 10.3.14.169:6800/1106 70 : [ERR] 1.b log bound mismatch, info (174 '4,419'34] actual [3'1,419'34]			

#### Associated revisions

Revision 9323f25a - 10/23/2011 03:41 AM - Sage Weil

osd: fix PG::Log::copy\_after wrt backlogs (again)

Commit 68fe748fc2d703623050e8f2a448a0fd31ca8a0f fixed half of this problem, but set this->tail incorrectly. If we read olog.tail, the entry we are on is a backlog entry, and probably not other.tail. Do not reset tail in this case because we already set it to other.tail above.

OTOH if we hit v, we do want to set this->tail to the current record as it is the one that precedes the first log entry.

This fixes an incorrect log.tail send to other nodes, which eventually propagates as a log bound mismatch. For example,

```
2011-10-22 17:33:18.654693 7f8a2fefe700 osd.4 2788 pg[1.1f( v 1627'28 (1627'28,1627'28) n=2 ec=1 les/c 2763/2782 2788/2788/2788) [4,0] r=0
mlcod 0'0 !hml peering] merge_log log(578'5,1627'28] from osd.0 into log(1627'28,1627'28]
2011-10-22 17:33:18.654706 7f8a2fefe700 osd.4 2788 pg[1.1f( v 1627'28 (1627'28,1627'28) n=2 ec=1 les/c 2763/2782 2788/2788/2788) [4,0] r=0
mlcod 0'0 !hml peering] merge_log extending tail to 578'5
2011-10-22 17:33:18.654720 7f8a2fefe700 osd.4 2788 pg[1.1f( v 1627'28 (578'5,1627'28) n=2 ec=1 les/c 2763/2782 2788/2788/2788) [4,0] r=0 (log
bound mismatch, empty) mlcod 0'0 !hml peering] merge_log result log(578'5,1627'28] missing(0) changed=1
```

This might fix #1526.

Signed-off-by: Sage Weil <[sage.weil@dreamhost.com](mailto:sage.weil@dreamhost.com)>

## History

---

### #1 - 09/09/2011 10:04 PM - Sage Weil

- translation missing: en.field\_position set to 46

### #2 - 09/19/2011 11:54 AM - Josh Durgin

Reproduced this with thrashing and rados bench - logs with full debugging are in vit:~joshd/thrash\_bound\_mismatch.

### #3 - 09/21/2011 03:26 PM - Sage Weil

- Assignee set to Sage Weil

### #4 - 09/21/2011 04:30 PM - Sage Weil

- Status changed from New to Resolved

should be fixed by [68fe748fc2d703623050e8f2a448a0fd31ca8a0f](https://github.com/sage/weil/commit/68fe748fc2d703623050e8f2a448a0fd31ca8a0f)

### #5 - 09/30/2011 02:43 PM - Josh Durgin

- Status changed from Resolved to In Progress

- Target version changed from v0.36 to v0.37

This happened again today. From teuthology:~teuthworker/archive/nightly\_coverage\_2011-09-30/903/ceph.log:

```
2011-09-30 11:30:50.310628 osd.14 10.3.14.189:6800/2117 27 : [ERR] 1.34 log bound mismatch, empty but (8'1,664'61]
2011-09-30 11:42:16.956732 osd.14 10.3.14.189:6800/2117 28 : [ERR] 1.34 log bound mismatch, empty but (8'1,664'61]
2011-09-30 11:42:34.729665 osd.14 10.3.14.189:6800/2117 29 : [ERR] 1.34 log bound mismatch, empty but (8'1,664'61]
2011-09-30 11:42:57.784552 osd.14 10.3.14.189:6800/2117 30 : [ERR] 1.34 log bound mismatch, empty but (8'1,664'61]
2011-09-30 12:47:31.721838 osd.7 10.3.14.208:6800/2088 20 : [ERR] 1.27 log bound mismatch, empty but (8'1,3005'1102]
2011-09-30 12:47:50.189615 osd.1 10.3.14.216:6800/2347 13 : [ERR] 1.27 log bound mismatch, info (8'1,3005'1102]+backlog actual [8'1,8'1]
2011-09-30 12:47:49.307363 osd.7 10.3.14.208:6800/2088 21 : [ERR] 1.27 log bound mismatch, info (8'1,3005'1102]+backlog actual [8'1,8'1]
2011-09-30 12:47:57.926531 osd.7 10.3.14.208:6800/2088 22 : [ERR] 1.27 log bound mismatch, info (8'1,3005'1102]+backlog actual [8'1,8'1]
2011-09-30 12:48:30.674807 osd.7 10.3.14.208:6800/2088 23 : [ERR] 1.27 log bound mismatch, empty but (8'1,3005'1102]
2011-09-30 12:48:42.504152 osd.7 10.3.14.208:6800/2088 24 : [ERR] 1.27 log bound mismatch, info (8'1,3005'1102]+backlog actual [8'1,8'1]
2011-09-30 12:48:43.006341 osd.1 10.3.14.216:6800/2347 14 : [ERR] 1.27 log bound mismatch, info (8'1,3005'1102]+backlog actual [8'1,8'1]
2011-09-30 12:49:20.090688 osd.7 10.3.14.208:6800/2088 25 : [ERR] 1.27 log bound mismatch, info (8'1,3005'1102]+backlog actual [8'1,8'1]
```

**#6 - 10/09/2011 08:39 PM - Sage Weil**

- Target version changed from v0.37 to v0.38

**#7 - 10/13/2011 10:12 AM - Sage Weil**

- Assignee deleted (Sage Weil)

**#8 - 10/19/2011 03:15 PM - Yehuda Sadeh**

another occurrence, running on swab. This may have led to pg version reset.

```
2011-10-19 14:12:49.013238 7f01e4471720 log [ERR] : 6.2 log bound mismatch, info (0'0,0'0] actual [7'1,27'2]
2011-10-19 14:12:49.013490 7f01e4471720 log [ERR] : 11.5 log bound mismatch, empty but (0'0,14'2]
2011-10-19 14:12:49.013651 7f01e4471720 log [ERR] : 6.6 log bound mismatch, info (0'0,9'2] actual [12'1,12'2]
2011-10-19 14:12:49.025347 7f01e4471720 log [ERR] : 11.3 log bound mismatch, empty but (0'0,14'10]
2011-10-19 14:12:49.067002 7f01e4471720 log [ERR] : 7.4 log bound mismatch, info (0'0,0'0] actual [8'1,8'2]
2011-10-19 14:12:49.067453 7f01e4471720 osd.0 47 pg[5.0( v 46'440 (0'0,46'440) n=3 ec=6 les/c 45/47 44/44/44)
```

```
[ ] r=0 (info mismatch, log(0'0,0'0)) (log bound mismatch, actual=[12'1,17'385]) lcod 0'0 mlcod 0'0 inactive]
read_log got dup 46'385 (last was 17'385, dropping that one)
2011-10-19 14:12:49.067466 7f01e4471720 log [ERR] : 5.0 read_log got dup 46'385 after 17'385
2011-10-19 14:12:49.070684 7f01e4471720 log [ERR] : 10.3 log bound mismatch, info (0'0,14'5] actual [11'1,42'1
46]
2011-10-19 14:12:49.070960 7f01e4471720 log [ERR] : 7.3 log bound mismatch, info (0'0,0'0) actual [8'1,27'2]
2011-10-19 14:12:49.077550 7f01e4471720 log [ERR] : 7.5 log bound mismatch, info (0'0,0'0) actual [8'1,27'2]
2011-10-19 14:12:49.080610 7f01e4471720 log [ERR] : 10.6 log bound mismatch, info (0'0,14'28] actual [11'1,42'
613]
2011-10-19 14:12:49.080745 7f01e4471720 log [ERR] : 11.7 log bound mismatch, empty but (0'0,14'31]
2011-10-19 14:12:49.115480 7f01e4471720 log [ERR] : 10.2 log bound mismatch, info (0'0,0'0) actual [11'1,42'63
7]
2011-10-19 14:12:49.140578 7f01e4471720 log [ERR] : 5.1 log bound mismatch, info (0'0,0'0) actual [7'1,7'4]
2011-10-19 14:12:49.176948 7f01e4471720 log [ERR] : 10.5 log bound mismatch, info (0'0,0'0) actual [11'1,42'63
2]
2011-10-19 14:12:49.209136 7f01e4471720 log [ERR] : 10.7 log bound mismatch, info (0'0,14'8] actual [11'1,42'6
19]
2011-10-19 14:12:49.246394 7f01e4471720 log [ERR] : 11.6 log bound mismatch, empty but (0'0,14'7]
2011-10-19 14:12:49.250584 7f01e4471720 log [ERR] : 6.3 log bound mismatch, info (0'0,0'0) actual [11'1,11'3]
2011-10-19 14:12:49.258714 7f01e4471720 log [ERR] : 5.3 log bound mismatch, info (0'0,0'0) actual [8'1,8'2]
2011-10-19 14:12:49.258983 7f01e4471720 log [ERR] : 11.0 log bound mismatch, info (0'0,0'0) actual [12'1,12'2]
```

**#9 - 10/25/2011 10:31 AM - Sage Weil**

- Status changed from In Progress to Resolved

- Assignee set to Sage Weil