

Ceph - Bug #1376

errant scrub stat mismatch logs after upgrade

08/08/2011 03:50 PM - John Leach

Status: Resolved	% Done: 0%
Priority: Normal	Spent time: 0.00 hour
Assignee:	
Category: OSD	
Target version: v0.35	
Source:	Reviewed:
Tags:	Affected Versions:
Backport:	ceph-qa-suite:
Regression: No	Pull request ID:
Severity: 3 - minor	Crash signature:
Description	
<p>upgraded from git commit #394537092d to git commit #68cbbf42c42, and after restarting the cluster I immediately saw many "scrub stat mismatch" errors:</p>	
<pre>2011-08-08 21:45:00.590610 log 2011-08-08 21:44:53.557073 osd1 10.42.77.90:6800/10375 1 : [ERR] 0.2 scrub stat mismatch, got 63/63 objects, 0/0 clones, 542633/18446744073701705641 bytes, 566/184 46744073709543990 kb. 2011-08-08 21:45:01.646482 log 2011-08-08 21:44:56.610022 osd2 10.219.16.42:6801/16231 1 : [ERR] 0.6 scrub stat mismatch, got 76/76 objects, 0/0 clones, 3822237/18446744073709179549 bytes, 3772/ 18446744073709551292 kb.2011-08-08 21:45:03.779548 log 2011-08-08 21:44:55.965413 osd3 10.200.35 .118:6800/1574 10 : [ERR] 0.15 scrub stat mismatch, got 83/83 objects, 0/0 clones, 2329158/1844674 4073707686470 bytes, 2321/18446744073709549841 kb.</pre>	
<p>as these came in from multi osds, and just after an upgrade (never seen them before and have done many upgrades), it looks more like a bug than real data corruption.</p>	
<p>Cluster layout is:</p>	
<pre>2011-08-08 22:45:29.712106 pg v35386: 800 pgs: 800 active+clean; 2679 MB data, 13032 MB used, 2 979 GB / 3152 GB avail 2011-08-08 22:45:29.713446 mds e91021: 2/2/2 up {0=1=up:active,1=0=up:active} 2011-08-08 22:45:29.713474 osd e102: 4 osds: 4 up, 4 in 2011-08-08 22:45:29.713520 log 2011-08-08 22:38:41.736291 osd3 10.200.35.118:6800/2525 297 : [ER R] 0.15 scrub 1 errors 2011-08-08 22:45:29.713583 mon e1: 3 mons at {0=10.126.174.94:6789/0,1=10.82.103.194:6789/0,2=10 .115.202.218:6789/0}</pre>	
<p>cluster is just a test cluster, no real data and at the time of the upgrade, had no clients accessing it.</p>	
<p>attached log of one of the osds after manually requesting a scrub (debug level 20)</p>	
Related issues:	
Related to Ceph - Bug #1453: osd: warn on object_info_t::size != st_size when...	Resolved 08/28/2011

History

#1 - 08/18/2011 03:28 PM - Sage Weil

- Target version set to v0.35

#2 - 08/20/2011 11:06 AM - Sage Weil

- translation missing: en.field_position set to 25

#3 - 08/20/2011 12:33 PM - John Leach

Just tried writing some data to the ceph filesystem on this cluster and got this message:

```
2011-08-20 19:26:24.661143 log 2011-08-20 19:16:14.568350 mds1 10.234.213.118:6800/3918 2 : [ERR] dir 200000
00441.200000000441 object missing on disk; some files may be lost
```

not sure if it's related in any way - never seen a message like that before.

#4 - 08/25/2011 03:44 PM - Greg Farnum

Missing objects on disk sure make it look like data corruption. Your cluster's pretty old, right? Is it still in this state?

#5 - 08/26/2011 02:14 PM - John Leach

it's a few weeks old yes, but there was no other evidence of of corruption (such as filesystem corruption).

I just deleted the osd data directory on osd1, re-added it to the cluster and let it rebuild and then ran a scrub and the errors came up again.

e.g:

```
2011-08-26 20:59:12.580635 log 2011-08-26 20:59:05.796145 osd1 10.42.77.90:6800/10085 20 : [ERR] 0.7 scrub s
tat mismatch, got 1769/1769 objects, 0/0 clones, 5458110067/5449721459 bytes, 5330431/5322239 kb.
2011-08-26 20:59:12.580635 log 2011-08-26 20:59:05.796161 osd1 10.42.77.90:6800/10085 21 : [ERR] 0.7 scrub 1
errors
```

This is with git commit 9538e87e0 now.

#6 - 08/28/2011 09:43 PM - Sage Weil

I think this is caused by an old bug. scrub needs to be fixed to properly detect (and ideally repair) it. See [#1453](#).

#7 - 08/29/2011 04:26 PM - Josh Durgin

- Status changed from New to 4

If you still have this cluster around, could you try applying [8293dfabb554883a30af549447995390fafa1f62](#) to see whether the problem is the old bug?

#8 - 09/05/2011 03:03 AM - John Leach

I upgraded to get that patch, but also got the on disk filestore update patch which was buggy and broke all my osds, so I can't test this any more, sorry.

#9 - 09/06/2011 09:49 PM - Sage Weil

- Target version changed from v0.35 to v0.36

#10 - 09/06/2011 10:07 PM - Sage Weil

- Status changed from 4 to Resolved

Ok. Well we're pretty sure what the inconsistency was, and we now complain about it (tho we don't repair it just yet). Making repair work is another bug, [#1474](#).

#11 - 09/06/2011 10:30 PM - Sage Weil

- Target version changed from v0.36 to v0.35

- translation missing: en.field_position deleted (69)

- translation missing: en.field_position set to 1

- translation missing: en.field_position changed from 1 to 898

Files

osd.3.log.gz	773 KB	08/08/2011	John Leach
--------------	--------	------------	------------