

Ceph - Bug #13962

"shard missing" errors in logs during Teuthology run

12/02/2015 06:35 PM - Piotr Dalek

Status:	Resolved	Start date:	12/02/2015
Priority:	Urgent	Due date:	
Assignee:	Sage Weil	% Done:	0%
Category:	OSD	Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:	Q/A	Reviewed:	
Tags:		Affected Versions:	v10.0.0
Backport:	infernalis	ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor		
Description			
In logs we can find:			
"2015-11-28 14:28:00.302371 osd.3 149.202.179.212:6804/9517 48 : cluster [ERR] 1.16 shard 2 missing 1/638ed856/target179214.teuthology10504-303/head" in cluster log			
Saw on http://149.202.162.14:8081/ubuntu-2015-11-27_20:07:13-rados:thrash-bp-delayed-pglog-index-v2---basic-openstack/13/ and http://pulpito.ovh.sepia.ceph.com:8081/teuthology-2015-11-29_21:00:02-rados-infernalis-distro-basic-openstack/22514/ (both logs attached).			
Related issues:			
Copied to Ceph - Backport #13979: infernalis : "shard missing" errors in logs...			Resolved

Associated revisions

Revision fb120d7b - 12/02/2015 07:50 PM - Sage Weil

osd: call on_new_interval on newly split child PG

We must call on_new_interval() on any interval change **and** on the creation of the PG. Currently we call it from PG::init() and PG::start_peering_interval(). However, PG::split_into() did not do so for the child PG, which meant that the new child feature bits were not properly initialized and the bitwise/nibblewise debug bit was not correctly set. That, in turn, could lead to various misbehaviors, the most obvious of which is scrub errors due to the sort order mismatch.

Fixes: #13962

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

Revision 7ac5b151 - 12/08/2015 08:17 AM - Sage Weil

osd: call on_new_interval on newly split child PG

We must call on_new_interval() on any interval change **and** on the creation of the PG. Currently we call it from PG::init() and

PG::start_peering_interval(). However, PG::split_into() did not do so for the child PG, which meant that the new child feature bits were not properly initialized and the bitwise/nibblewise debug bit was not correctly set. That, in turn, could lead to various misbehaviors, the most obvious of which is scrub errors due to the sort order mismatch.

Fixes: #13962

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

(cherry picked from commit fb120d7b2da5715e7f7d1baa65bfa70d2e5d807a)

History

#1 - 12/02/2015 07:04 PM - Sage Weil

- Status changed from New to In Progress
- Assignee set to Sage Weil
- Source changed from other to Q/A

/a/sage-2015-11-30_05:29:21-rados-wip-sage-testing---basic-multi/1164042

#2 - 12/02/2015 07:24 PM - Sage Weil

nibblewise vs bitwise mismatch:

```
2015-11-30 16:59:47.940540 7fe9c9839700 10 osd.1 pg_epoch: 81 pg[1.21( v 81'299 (65'213,81'299) local-les=62 n
=85 ec=5 les/c/f 62/65/0 56/58/49) [4,1] r=1 lpr=62 pi=49-57/2 luod=0'0 crt=74'297 lcod 76'298 active] be_scan
_list scanning 42 objects
2015-11-30 16:59:47.933462 7ff4e0cd4700 10 osd.4 pg_epoch: 81 pg[1.21( v 81'299 (65'213,81'299) local-les=62 n
=85 ec=5 les/c/f 62/65/0 56/58/49) [4,1] r=0 lpr=58 crt=74'297 lcod 76'298 mlcod 76'298 active+clean+scrubbing
NIBBLEWISE] be_scan_list scanning 24 objects
```

#3 - 12/02/2015 07:54 PM - Sage Weil

- Status changed from In Progress to Need Review

<https://github.com/ceph/ceph/pull/6778>

#4 - 12/04/2015 12:38 AM - Sage Weil

- Status changed from Need Review to Pending Backport
- Backport set to infernalis

#5 - 12/04/2015 06:54 AM - Abhishek Varshney

- Copied to Backport #13979: infernalis : "shard missing" errors in logs during Teuthology run added

#6 - 02/08/2016 03:33 PM - Nathan Cutler

- Status changed from Pending Backport to Resolved

Files

teuthology-logs.zip

666 KB

12/02/2015

Piotr Dalek