Ceph - Bug #629

cosd segfaults when deleting a pool containing degraded objects

12/02/2010 04:54 PM - John Leach

Status:	Resolved	**% Done:**	0%
Priority:	Normal	**Spent time:**	0.00 hour
Assignee:	Colin McCabe		
Category:			
Target version:	v0.25		
Source:		**Reviewed:**	
Tags:		**Affected Versions:**	
Backport:		**ceph-qa-suite:**	
Regression:	No	**Pull request ID:**	
Severity:	3 - minor	**Crash signature:**	

Description

started a 4 node osd cluster. created some pools with some objects in them. killed one osd node. waited for it to be noticed and cluster to become degraded. deleted 3 pools containing degraded objects (using rados rmpool) and shortly afterward, other cosd processes segfault:

```
2010-12-03 00:48:39.443120 7fffeab28710 osd1 5181 pg[385.0( v 1158'1219 lc 0'0 (1158'1217,1158'121
9]+backlog n=1219 ec=1155 les=5138 5158/5158/5158) [] r=-1 (info mismatch, log(0'0,0'0]) stray DEL
ETING] write_log to 0~0
2010-12-03 00:48:39.443155 7fffeab28710 osd1 5181 _remove_pg 385.0 0 objects
2010-12-03 00:48:39.443163 7fffeab28710 osd1 5181 _remove_pg 385.0 flushing store
2010-12-03 00:48:39.443815 7fffeab28710 osd1 5181 _remove_pg 385.0 taking osd_lock
2010-12-03 00:48:39.443832 7fffeab28710 osd1 5181 _remove_pg 385.0 removing final

Program received signal SIGSEGV, Segmentation fault.
[Switching to Thread 0x7fffeab28710 (LWP 13457)]
0x00000000004c37b4 in OSD::_put_pool(int) ()

(gdb) bt

#0  0x00000000004c37b4 in OSD::_put_pool(int) ()
#1  0x00000000004d6e7a in OSD::_remove_pg(PG*) ()
#2  0x00000000005d0a4f in ThreadPool::worker() ()
#3  0x00000000004feeed in ThreadPool::WorkThread::entry() ()
#4  0x0000000000470baa in Thread::_entry_func(void*) ()
#5  0x00007ffff79c29ca in start_thread () from /lib/libpthread.so.0
#6  0x00007ffff694070d in clone () from /lib/libc.so.6
#7  0x0000000000000000 in ?? ()
```

full all-thread backtrace attached.

Related issues:

Related to Ceph - Bug #696: osd: _put_pool, assert(p->num_pg > 0)	**Resolved**	**01/09/2011**

History

#1 - 12/02/2010 05:02 PM - Colin McCabe

Looks like some kind of lifecycle issue related to deleting pools.

OSD::_remove_pg does a _put_pool, and that does a _lookup_pool. That _lookup_pool must be returning NULL-- I think that is the only way to get a segfault in OSD::_put_pool.

#2 - 12/06/2010 10:51 AM - Sage Weil

- Target version set to v0.25

#3 - 12/16/2010 04:24 PM - Sage Weil

- Assignee set to Colin McCabe

#4 - 12/17/2010 01:29 PM - Colin McCabe

- Status changed from New to 7

This shouldn't happen again c3a24fc5d31d53e3db911be900b9067584f0e07e

It still might be interesting to see the logs leading up to the original crash, though. Post them if you have 'em!

#5 - 01/06/2011 02:54 PM - Sage Weil

- Status changed from 7 to Resolved

Files

gdb.txt	15.1 KB	12/02/2010	John Leach