

Ceph - Bug #579

OSD::sched_scrub: FAILED assert(pg_map.count(pgid))

11/15/2010 10:51 AM - Colin McCabe

Status:	Resolved	Start date:	11/15/2010
Priority:	Normal	Due date:	
Assignee:	Sage Weil	% Done:	0%
Category:	OSD	Estimated time:	0.00 hour
Target version:	v0.24	Spent time:	0.50 hour
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor		

Description

On unfound_last_epoch_clean at commit commit:7201497f2feef6a2bbd0baf89e3a14b8a880e79f

I found this assert when running

```
./test/test_unfound.sh stray_test
```

At first glance, seems to be a refcount issue in the scrub code.

=====

```
osd/OSD.cc: In function 'PG* OSD::_lookup_lock_pg(pg_t)':
osd/OSD.cc:956: FAILED assert(pg_map.count(pgid))
ceph version 0.24~rc (commit:7201497f2feef6a2bbd0baf89e3a14b8a880e79f)
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x34) [0x81ee4e]
2: (OSD::_lookup_lock_pg(pg_t)+0x6f) [0x6d3d57]
3: (OSD::sched_scrub()+0x2e9) [0x6e4445]
4: (OSD::tick()+0x204) [0x6f168e]
5: (OSD::C_Tick::finish(int)+0x1c) [0x7613bc]
6: (SafeTimer::timer_thread()+0x189) [0x81bcf5]
7: (SafeTimerThread::entry()+0x19) [0x81dd73]
8: (Thread::_entry_func(void*)+0x20) [0x66496a]
9: (()+0x68ba) [0x7fb807d118ba]
10: (clone()+0x6d) [0x7fb806a7002d]
```

NOTE: a copy of the executable, or `objdump -rDS <executable>` is needed to interpret this.

- Caught signal (ABRT) **
ceph version 0.24~rc (commit:7201497f2feef6a2bbd0baf89e3a14b8a880e79f)
1: (ceph::BackTrace::BackTrace(int)+0x2a) [0x81f134]
2: (sigabrt_handler(int)+0x41) [0x8302eb]
3: (()+0x321e0) [0x7fb8069d31e0]
4: (gsignal()+0x35) [0x7fb8069d3165]
5: (abort()+0x180) [0x7fb8069d5f70]
6: (_gnu_cxx::verbose_terminate_handler()+0x115) [0x7fb807266dc5]
7: (()+0xcb166) [0x7fb807265166]
8: (()+0xcb193) [0x7fb807265193]
9: (()+0xcb28e) [0x7fb80726528e]
10: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x217) [0x81f031]
11: (OSD::_lookup_lock_pg(pg_t)+0x6f) [0x6d3d57]
12: (OSD::sched_scrub()+0x2e9) [0x6e4445]
13: (OSD::tick()+0x204) [0x6f168e]
14: (OSD::C_Tick::finish(int)+0x1c) [0x7613bc]
15: (SafeTimer::timer_thread()+0x189) [0x81bcf5]
16: (SafeTimerThread::entry()+0x19) [0x81dd73]
17: (Thread::_entry_func(void*)+0x20) [0x66496a]

18: (()+0x68ba) [0x7fb807d118ba]
19: (clone()+0x6d) [0x7fb806a7002d]

Related issues:

Related to Ceph - Bug #612: OSD: Crash during auto scrub

Resolved

11/24/2010

History

#1 - 11/15/2010 10:56 AM - Colin McCabe

Some more information about this bug.

OSD1 and OSD2 have a PG named 0.6
OSD0 does not.

```
=====  
cmccabe@flab:~/src/ceph/src$ find dev | grep 0.6  
dev/osd2/current/meta/pglog_0.6_0  
dev/osd2/current/meta/pginfo_0.6_0  
dev/osd2/current/0.6_head  
dev/osd2/current/0.6_head/obj07_head  
dev/osd1/current/meta/pglog_0.6_0  
dev/osd1/current/meta/pginfo_0.6_0  
dev/osd1/current/0.6_head  
dev/osd1/current/0.6_head/obj07_head =====
```

The last thing osd0 printed out before it crashed was

```
2010-11-15 10:44:13.702856 7fb80456d710 osd0 12 on 0.000000 0.6
```

So I believe that after the test ran this line:

```
./ceph osd in 0
```

The crash was only a matter of time because the OSD simply did not have the PGs that the other OSDs had. It had been out for a while.

#2 - 11/15/2010 11:19 AM - Sage Weil

- Category set to OSD
- Status changed from New to Resolved
- Assignee set to Sage Weil

- Target version set to v0.24

commit:f46f674261bf65a6f7f6313fb688ec4773f526b5