

## Ceph - Bug #872

### osd: crash due to missing pginfo

03/10/2011 05:52 AM - Wido den Hollander

<b>Status:</b>	Resolved	<b>Start date:</b>	03/10/2011
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Sage Weil	<b>% Done:</b>	0%
<b>Category:</b>	OSD	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>		<b>Spent time:</b>	2.00 hours
<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>	

#### Description

I just upgraded "noisy" and saw osd1 go down after restart with:

```
2011-03-10 14:20:29.799991 7f717c26d720 filestore(/var/lib/ceph/osd.1) collection_getattr /var/lib/ceph/osd.1/current/3.7c9_head 'info'
2011-03-10 14:20:29.800049 7f717c26d720 filestore(/var/lib/ceph/osd.1) collection_getattr /var/lib/ceph/osd.1/current/3.7c9_head 'info' = 309
2011-03-10 14:20:29.800061 7f717c26d720 filestore(/var/lib/ceph/osd.1) read /var/lib/ceph/osd.1/current/meta/pginfo_3.7c9_0 0~0
2011-03-10 14:20:29.800102 7f717c26d720 filestore(/var/lib/ceph/osd.1) FileStore::read(/var/lib/ceph/osd.1/current/meta/pginfo_3.7c9_0): open error error 2: No such file or directory
*** Caught signal (Aborted) **
in thread 0x7f717c26d720
ceph version 0.26~rc (commit:1f120284ed80ee1258b556fbedacab209098a0d1)
1: /usr/bin/cosd() [0x61b078]
2: (()+0xf8f0) [0x7f717bc4e8f0]
3: (gsignal()+0x35) [0x7f717a81ea75]
4: (abort()+0x180) [0x7f717a8225c0]
5: (__gnu_cxx::__verbose_terminate_handler()+0x115) [0x7f717b0d48e5]
6: (()+0xcad16) [0x7f717b0d2d16]
7: (()+0xcad43) [0x7f717b0d2d43]
8: (()+0xcae3e) [0x7f717b0d2e3e]
9: (ceph::buffer::list::iterator::copy(unsigned int, char*)+0x12c) [0x468dbc]
10: (void decode<unsigned int, PG::Interval>(std::map<unsigned int, PG::Interval, std::less<unsigned int>, std::allocator<std::pair<unsigned int const, PG::Interval> > >&, ceph::buffer::list::iterator&)+0x31) [0x577b81]
11: (PG::read_state(ObjectStore*)+0x32b) [0x55e78b]
12: (OSD::load_pgs()+0x1b4) [0x4fbfb4]
13: (OSD::init()+0x517) [0x519f87]
14: (main()+0x1770) [0x4662f0]
15: (__libc_start_main()+0xfd) [0x7f717a809c4d]
16: /usr/bin/cosd() [0x4646e9]
```

**3.7c9** was one of the PG's which kept blocking ([#847](#))

A search for the pg info brought me:

```
root@noisy:/var/log/ceph# find /var/lib/ceph/ -name pginfo_3.7c9_0
/var/lib/ceph/osd.1/current.remove.me.846930886/meta/pginfo_3.7c9_0
/var/lib/ceph/osd.1/snap_856539/meta/pginfo_3.7c9_0
root@noisy:/var/log/ceph#
```

## History

---

### #1 - 03/10/2011 10:26 AM - Sage Weil

- Assignee set to Sage Weil

### #2 - 03/10/2011 11:43 AM - Sage Weil

Ah, this is my fault. I made a copy of the files in 3.7c9 in a subdir called 't' (they were missing xattrs... :/) while debugging the old issue. And then when cosd went and removed all objects, the rmdir on 3.7c9\_head failed (not empty). And now when it starts up it sees the dir but no info, and crashes. Just remove the dir from the most recent snap\_\* dir and it should start right up.

I'm not sure what the proper behavior here should be. We can throw the error when the rmdir fails and crash then? Or log something and continue?

### #3 - 03/10/2011 11:56 AM - Wido den Hollander

I do not think that crashing due to one faulty dir is what I'd do, but on the other hand, it will force an admin to keep the OSD's datadir 'sane', which might prevent further issues in the future.

I could live with both options, you could make it a config option which defaults to crashing?

### #4 - 03/10/2011 12:22 PM - Wido den Hollander

Just thought about this, will this something an admin would run into? I ran into this due to the recovery issue. But in a real production env, wouldn't you just wipe the OSD? It's not likely to come back, unless someone messes around with the OSD datadir.

### #5 - 03/10/2011 01:11 PM - Sage Weil

Wido den Hollander wrote:

Just thought about this, will this something an admin would run into? I ran into this due to the recovery issue. But in a real production env, wouldn't you just wipe the OSD? It's not likely to come back, unless someone messes around with the OSD datadir.

Right. This only happened because I polluted things with files that shouldn't be there.

For now I'll just assert on ENOTEMPTY; that makes the most sense given the existing error handling.

### #6 - 03/10/2011 01:13 PM - Sage Weil

- Status changed from New to Resolved