

## Ceph - Bug #12162

### pg\_interval\_t::check\_new\_interval - for ec pool, should not rely on min\_size to determine if the PG was active at the interval

06/25/2015 07:05 PM - Guang Yang

<b>Status:</b> Resolved	<b>% Done:</b> 0%
<b>Priority:</b> Urgent	<b>Spent time:</b> 0.00 hour
<b>Assignee:</b> Guang Yang	
<b>Category:</b>	
<b>Target version:</b>	
<b>Source:</b> other	<b>Affected Versions:</b>
<b>Tags:</b>	<b>ceph-qa-suite:</b>
<b>Backport:</b> hammer, firefly	<b>Pull request ID:</b>
<b>Regression:</b> No	<b>Crash signature (v1):</b>
<b>Severity:</b> 3 - minor	<b>Crash signature (v2):</b>
<b>Reviewed:</b>	
<b>Description</b>	
One PG on our cluster stuck at peering+down forever, log shows the peering was blocked by an out/down OSD	
<pre>2015-06-23 21:27:59.948809 7f26d2bd9700 10 osd.52 pg_epoch: 27468 pg[3.1e3fs0( v 25576'94308 (1525 1'84307,25576'94308] local-les=25568 n=73284 ec=1152 les/c 25568/18381 27457/27457/27457) [52,23,1 0,456,433,200,388,330,493,104,426] r=0 lpr=27457 pi=18380-27456/559 crt=25298'94302 lcod 0'0 mlcod 0'0 peering] PriorSet: build_prior final: probe 10(2),22(1),23(1),24(2),46(0),52(0),104(9),200(5 ),249(3),254(10),265(6),330(7),388(6),426(10),433(4),450(7),456(3),493(8) down 243 blocked_by {243 =0} pg_down</pre>	
The actual problem came from that when building the PriorSet, it blindly used the pool's min_size when check if the PG was r/w during the interval:	
<pre>2015-06-23 21:28:00.357787 7f26d13d6700 10 osd.52 pg_epoch: 27471 pg[3.1e3fs0( v 25576'94308 (1525 1'84307,25576'94308] local-les=25568 n=73284 ec=1152 les/c 25568/18381 27471/27471/27471) [52,23,1 0,456,433,200,388,330,493,104,426] r=0 lpr=27471 pi=18380-27470/561 crt=25298'94302 lcod 0'0 mlcod 0'0 peering] PriorSet: build_prior interval(25614-25615 up [2147483647,2147483647,2147483647,214 7483647,2147483647,200,2147483647,2147483647,243,104,426] (200) acting [2147483647,2147483647,21474 83647,2147483647,2147483647,200,2147483647,2147483647,243,104,426] (200) maybe_went_rw)</pre>	
Ceph version: v0.80.4	
Credit goes to Sam for the analysis, thanks Sam!	
<b>Related issues:</b>	
Copied to Ceph - Backport #12488: pg_interval_t::check_new_interval - for ec ...	<b>Rejected</b>
Copied to Ceph - Backport #12489: pg_interval_t::check_new_interval - for ec ...	<b>Resolved</b> <b>06/25/2015</b>

### Associated revisions

#### Revision 68492744 - 07/13/2015 05:42 PM - Guang G Yang

osd: pg\_interval\_t::check\_new\_interval should not rely on pool.min\_size to determine if the PG was active

If the pool's min\_size is set improperly, during peering, pg\_interval\_t::check\_new\_interval might wrongly determine the PG's state and cause the PG to stuck at down+peering forever

Fixes: #12162

Signed-off-by: Guang Yang [yguang@yahoo-inc.com](mailto:yguang@yahoo-inc.com)

### Revision cd11b887 - 07/28/2015 01:11 PM - Guang G Yang

osd: pg\_interval\_t::check\_new\_interval should not rely on pool.min\_size to determine if the PG was active

If the pool's min\_size is set improperly, during peering, pg\_interval\_t::check\_new\_interval might wrongly determine the PG's state and cause the PG to stuck at down+peering forever

Fixes: #12162

Signed-off-by: Guang Yang [yguang@yahoo-inc.com](mailto:yguang@yahoo-inc.com)

(cherry picked from commit 684927442d81ea08f95878a8af69d08d3a14d973)

Conflicts:

src/osd/PG.cc

because PG::start\_peering\_interval has an assert that is not found in hammer in the context

src/test/osd/types.cc

because include/stringify.h is not included by types.cc in hammer

## History

---

### #1 - 06/25/2015 09:55 PM - Guang Yang

- Subject changed from *PG is stuck at down+peering forever* to *pg\_interval\_t::check\_new\_interval - for ec pool, should not relay on min\_size to determine if the PG was active at the interval*

### #2 - 06/26/2015 04:22 PM - Guang Yang

- Subject changed from *pg\_interval\_t::check\_new\_interval - for ec pool, should not relay on min\_size to determine if the PG was active at the interval* to *pg\_interval\_t::check\_new\_interval - for ec pool, should not rely on min\_size to determine if the PG was active at the interval*

### #3 - 06/26/2015 06:23 PM - Guang Yang

<https://github.com/ceph/ceph/pull/5093>

### #4 - 06/30/2015 02:39 PM - Samuel Just

- Priority changed from *Normal* to *Urgent*

### #5 - 07/14/2015 02:27 PM - Samuel Just

- Status changed from *New* to *7*

### #6 - 07/23/2015 05:49 PM - Samuel Just

- Status changed from *7* to *Pending Backport*

- Backport set to *hammer, firefly*

### #7 - 01/28/2016 12:19 PM - Loïc Dachary

- Status changed from *Pending Backport* to *Resolved*