

Ceph - Bug #19814

pg_create message at epoch 0 (on previously created pg)

05/01/2017 01:58 AM - Sage Weil

Status:	Resolved	Start date:	04/30/2017
Priority:	Immediate	Due date:	
Assignee:	Kefu Chai	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor		

Description

```
2017-04-30 00:45:40.451328 7fe1c36c7700 10 mon.b@0(leader).osd e416 scan_for_creating_pgs adding 1
95.0 at 416
...
2017-04-30 00:45:40.602606 7fe1c5ecc700 10 register_pg will create 195.0 primary 5 acting [5,1]
...
2017-04-30 00:45:40.635847 7fe1c1ec4700 10 mon.b@0(leader).osd e417 update_creating_pgs will instr
uct osd.5 to create 195.0
...
2017-04-30 00:45:40.635921 7fe1c1ec4700 1 -- 172.21.15.62:6789/0 --> 172.21.15.62:6811/21034 -- o
sd_pg_create(e417 195.0:417 195.1:417 195.2:417 195.4:417 195.6:417) v3 -- 0x564972344900 con 0
...
2017-04-30 00:45:40.727517 7fe1c5ecc700 20 mon.b@0(leader).pg v462 refreshing pg 195.0 0:0 creati
ng
...
2017-04-30 00:45:40.929513 7fe1c0ec2700 1 -- 172.21.15.62:6789/0 <== osd.5 172.21.15.62:6811/2103
4 29 ==== osd_pg_created(195.0) v1 ==== 35+0+0 (1037313915 0 0) 0x564973555400 con 0x564973593000
...
2017-04-30 00:45:42.028974 7fe1c0ec2700 20 mon.b@0(leader).osd e419 send_pg_creates will create 19
5.0 at 0
2017-04-30 00:45:42.028974 7fe1c16c3700 10 osdmap epoch 419 mapping took 0.000107 seconds
2017-04-30 00:45:42.028977 7fe1c0ec2700 20 mon.b@0(leader).osd e419 send_pg_creates will create 19
5.1 at 0
2017-04-30 00:45:42.028979 7fe1c0ec2700 20 mon.b@0(leader).osd e419 send_pg_creates will create 19
5.2 at 0
2017-04-30 00:45:42.028978 7fe1c16c3700 10 mon.b@0(leader).osd e419 update_creating_pgs will instr
uct osd.5 to create 195.0
2017-04-30 00:45:42.028980 7fe1c0ec2700 20 mon.b@0(leader).osd e419 send_pg_creates will create 19
5.4 at 0
2017-04-30 00:45:42.028982 7fe1c0ec2700 20 mon.b@0(leader).osd e419 send_pg_creates will create 19
5.6 at 0
2017-04-30 00:45:42.028982 7fe1c16c3700 10 mon.b@0(leader).osd e419 update_creating_pgs will instr
uct osd.5 to create 195.1
2017-04-30 00:45:42.028984 7fe1c16c3700 10 mon.b@0(leader).osd e419 update_creating_pgs will instr
uct osd.5 to create 195.2
2017-04-30 00:45:42.028985 7fe1c16c3700 10 mon.b@0(leader).osd e419 update_creating_pgs will instr
uct osd.0 to create 195.3
2017-04-30 00:45:42.028983 7fe1c0ec2700 1 -- 172.21.15.62:6789/0 --> 172.21.15.62:6811/21034 -- o
sd_pg_create(e418 195.0:0 195.1:0 195.2:0 195.4:0 195.6:0) v3 -- 0x564973617a80 con 0
```

/a/sage-2017-04-29_22:06:05-rados-wip-sage-testing2---basic-smithi/1082864

History

#1 - 05/01/2017 11:26 AM - Kefu Chai

- Assignee set to Kefu Chai

#2 - 05/01/2017 12:04 PM - Kefu Chai

seems creating_pgs_by_osd_epoch is not in sync with pg_creatings, the pg_creatings at that moment does not have pg 195.0, but when sending "osd_pg_create()", we look up the "created" epoch in pg_creatings, the default ctor is called for the "[]" operator. will take a look later.

#3 - 05/01/2017 03:54 PM - Kefu Chai

- Status changed from New to Need Review

<https://github.com/ceph/ceph/pull/14897>

#4 - 05/02/2017 01:31 AM - Sage Weil

- Status changed from Need Review to Resolved