

RADOS - Bug #23648

max-pg-per-osd.from-primary fails because of activating pg

04/11/2018 04:56 AM - Kefu Chai

Status: New	% Done: 0%
Priority: Normal	Spent time: 0.00 hour
Assignee: Kefu Chai	
Category: Correctness/Safety	
Target version:	
Source:	Affected Versions:
Tags:	ceph-qa-suite:
Backport:	Component(RADOS): Monitor
Regression: No	Pull request ID:
Severity: 3 - minor	Crash signature:
Reviewed:	

Description

the reason why we have activating pg when the number of pg is under the hard limit of max-pg-per-osd is that:

1. osd.1 received the osdmap instructing it to create pg 1.0, so it updated its replica osd.2 to create pg 1.0
2. osd.2 is capped by the max-pg-per-osd when it's about to create pg 1.0. so it dropped the request to create the pg.
3. the pool 15 gets removed in osdmap#58.
4. but the updated osdmap is not sent to osd.1 or osd.2 before the wait_for_clean() times out.

Related issues:

Related to RADOS - Bug #23610: pg stuck in activating because of dropped pg-t...	Resolved	04/10/2018
--	-----------------	-------------------

History

#1 - 04/11/2018 04:57 AM - Kefu Chai

- Related to Bug #23610: pg stuck in activating because of dropped pg-temp message added