

RADOS - Bug #46914

mon: stuck osd_pgtemp message forwards

08/12/2020 11:22 PM - Greg Farnum

Status: Resolved	% Done: 0%
Priority: Normal	Spent time: 0.00 hour
Assignee: Greg Farnum	
Category: Administration/Usability	
Target version:	
Source:	Affected Versions:
Tags:	ceph-qa-suite:
Backport: octopus, nautilus	Component(RADOS): Monitor
Regression: No	Pull request ID: 36593
Severity: 3 - minor	Crash signature (v1):
Reviewed:	Crash signature (v2):
Description	
https://bugzilla.redhat.com/show_bug.cgi?id=1866257	
We've seen osd_pgtemp messages which are forwarded to the leader not getting replies.	
Related issues:	
Copied to RADOS - Backport #47091: octopus: mon: stuck osd_pgtemp message for...	Resolved
Copied to RADOS - Backport #47092: nautilus: mon: stuck osd_pgtemp message fo...	Resolved

History

#1 - 08/12/2020 11:24 PM - Greg Farnum

Looking through the source, there's one clear way this happens: the leader may decide that a message can get dropped in preprocess_pgtemp() where the peon disagreed. That will usually involve some kind of race in changing states (eg, OSD getting marked down), but it seems possible.

So, just mark the op as no_reply when dropping the message.

#2 - 08/13/2020 12:28 AM - Greg Farnum

- Status changed from New to Fix Under Review

- Backport set to octopus, nautilus

<https://github.com/ceph/ceph/pull/36593>

#3 - 08/21/2020 01:53 AM - Kefu Chai

- Status changed from Fix Under Review to Pending Backport

#4 - 08/22/2020 07:44 PM - Nathan Cutler

- Copied to Backport #47091: octopus: mon: stuck osd_pgtemp message forwards added

#5 - 08/22/2020 07:44 PM - Nathan Cutler

- Copied to Backport #47092: nautilus: mon: stuck osd_pgtemp message forwards added

#6 - 09/15/2020 11:59 PM - Neha Ojha

- Pull request ID set to 36593

#7 - 09/30/2020 03:40 PM - Nathan Cutler

- *Status changed from Pending Backport to Resolved*

While running with --resolve-parent, the script "backport-create-issue" noticed that all backports of this issue are in status "Resolved" or "Rejected".