

Ceph - Bug #8815

mon: scrub error (osdmap encoding mismatch?) upgrading from 0.80 to ~0.80.2

07/11/2014 09:19 AM - Sage Weil

Status: Resolved	% Done: 0%
Priority: Urgent	Spent time: 0.00 hour
Assignee: Sage Weil	
Category:	
Target version:	
Source: Q/A	Affected Versions:
Tags:	ceph-qa-suite:
Backport: firefly	Pull request ID:
Regression: No	Crash signature (v1):
Severity: 3 - minor	Crash signature (v2):
Reviewed:	
Description	
ubuntu@teuthology:/a/sage-2014-07-10_20:55:57-upgrade:firefly-firefly-testing-basic-plana/353414	
description: upgrade/firefly/{0-cluster/start.yaml 1-install/v0.80.yaml 2-workload/testrados.yaml 3-upgrade-sequence/upgrade-mon-osd-mds.yaml 4-final/{monthrash.yaml osdthrash.yaml testrgw.yaml}}	
duration: 3798.3317320346832	
failure_reason: "'2014-07-11 02:35:54.764798 mon.0 10.214.133.24:6789/0 9 : [ERR] mon.0 ScrubResult(keys {auth=13,logm=1185,mdsmap=12,monmap=3,osdmap=1023,pgmap=516,pgmap_meta=6,pgmap_osd=6,pgmap_pg=72} crc {auth=4249049519,logm=1976909497,mdsmap=3432914802,monmap=132340688,osdmap=2754814452,pgmap=1350280295,pgmap_meta=1503530791,pgmap_osd=2157513224,pgmap_pg=290475359})'"	
in cluster log'	
Related issues:	
Copied to Ceph - Backport #12289: mon: scrub error (osdmap encoding mismatch?...	Resolved 07/11/2014

Associated revisions

Revision c93da05b - 07/11/2014 06:31 PM - Sage Weil

osd/osd_types: be pedantic about encoding last_force_op_resend without feature bit

The addition of the value is completely backward compatible, but if the mon feature bits don't match it can cause monitor scrub noise (due to the parallel OSDMap encoding). Avoid that by only adding the new field if the feature (which was added 2 patches after the encoding, see 3152faf79f498a723ae0fe44301ccb21b15a96ab and 45e79a17a932192995f8328ae9f6e8a2a6348d10.

Fixes: #8815

Backport: firefly

Signed-off-by: Sage Weil <sage@redhat.com>

Revision a8f3d6e1 - 07/10/2015 08:38 PM - Sage Weil

osd/osd_types: be pedantic about encoding last_force_op_resend without feature bit

The addition of the value is completely backward compatible, but if the mon feature bits don't match it can cause monitor scrub noise (due to the parallel OSDMap encoding). Avoid that by only adding the new field if the feature (which was added 2 patches after the encoding, see 3152faf79f498a723ae0fe44301ccb21b15a96ab and 45e79a17a932192995f8328ae9f6e8a2a6348d10).

Fixes: #8815

Backport: firefly

Signed-off-by: Sage Weil <sage@redhat.com>

(cherry picked from commit c93da05bd3dd6b73e8cfa182c6349359b042717c)

Conflicts:

src/osd/osd_types.cc

drop encode_compat variable which is otherwise unused in this context

History

#1 - 07/11/2014 09:42 AM - Sage Weil

- Priority changed from *Immediate* to *Urgent*

The problem is we added the pool_op_resend to pg_pool_t encoding, but didn't encoding conditionally based on teh feature. That is totally fine, the encodings are fully backward compatible, except that it trips up the mon scrub. Will fix it up in the next release.

#2 - 07/14/2014 03:08 PM - Sage Weil

- Status changed from *12* to *Resolved*

#3 - 07/10/2015 08:32 PM - Nathan Cutler

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

- Backport set to *firefly*

- Regression set to *No*

#4 - 09/10/2015 09:51 AM - Nathan Cutler

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