

RADOS - Bug #24612

FAILED assert(osdmap_manifest.pinned.empty()) in OSDMonitor::prune_init()

06/21/2018 06:48 PM - Neha Ojha

Status:	Resolved	Start date:	06/21/2018
Priority:	Urgent	Due date:	
Assignee:	Joao Eduardo Luis	% Done:	0%
Category:	Correctness/Safety	Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:	mimic	ceph-qa-suite:	
Regression:	No	Component(RADOS):	Monitor
Severity:	3 - minor	Pull request ID:	
Description			
<pre>2018-06-21T05:40:08.919 INFO:tasks.ceph.mon.b.smithi141.stderr:/build/ceph-14.0.0-586-g47c0cd8/src/mon/OSDMonitor.cc: In function 'void OSDMonitor::prune_init()' thread 7f9878853700 time 2018-06-21 05:40:08.918368 2018-06-21T05:40:08.919 INFO:tasks.ceph.mon.b.smithi141.stderr:/build/ceph-14.0.0-586-g47c0cd8/src/mon/OSDMonitor.cc: 1736: FAILED assert(osdmap_manifest.pinned.empty()) 2018-06-21T05:40:08.920 INFO:tasks.ceph.mon.b.smithi141.stderr: ceph version 14.0.0-586-g47c0cd8 (47c0cd87f07352f00bb964a25af79776e61a07c6) nautilus (dev) 2018-06-21T05:40:08.921 INFO:tasks.ceph.mon.b.smithi141.stderr: 1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x102) [0x7f9882bcc6b2] 2018-06-21T05:40:08.921 INFO:tasks.ceph.mon.b.smithi141.stderr: 2: (()+0x2e6877) [0x7f9882bcc877] 2018-06-21T05:40:08.922 INFO:tasks.ceph.mon.b.smithi141.stderr: 3: (OSDMonitor::prune_init()+0x1cd) [0x55bf478be41d] 2018-06-21T05:40:08.922 INFO:tasks.ceph.mon.b.smithi141.stderr: 4: (OSDMonitor::do_prune(std::shared_ptr<MonitorDBStore::Transaction>)+0x236) [0x55bf478e3266] 2018-06-21T05:40:08.922 INFO:tasks.ceph.mon.b.smithi141.stderr: 5: (OSDMonitor::encode_pending(std::shared_ptr<MonitorDBStore::Transaction>)+0x7b) [0x55bf478e440b] 2018-06-21T05:40:08.923 INFO:tasks.ceph.mon.b.smithi141.stderr: 6: (PaxosService::propose_pending()+0x113) [0x55bf478b18d3] 2018-06-21T05:40:08.923 INFO:tasks.ceph.mon.b.smithi141.stderr: 7: (C_MonContext::finish(int)+0x39) [0x55bf477528a9] 2018-06-21T05:40:08.923 INFO:tasks.ceph.mon.b.smithi141.stderr: 8: (Context::complete(int)+0x9) [0x55bf4778a9a9] 2018-06-21T05:40:08.923 INFO:tasks.ceph.mon.b.smithi141.stderr: 9: (SafeTimer::timer_thread()+0x18b) [0x7f9882bc902b] 2018-06-21T05:40:08.924 INFO:tasks.ceph.mon.b.smithi141.stderr: 10: (SafeTimerThread::entry()+0xd) [0x7f9882bca5ed] 2018-06-21T05:40:08.924 INFO:tasks.ceph.mon.b.smithi141.stderr: 11: (()+0x76ba) [0x7f98824c86ba] 2018-06-21T05:40:08.925 INFO:tasks.ceph.mon.b.smithi141.stderr: 12: (clone()+0x6d) [0x7f98811ec41d] 2018-06-21T05:40:08.925 INFO:tasks.ceph.mon.b.smithi141.stderr: NOTE: a copy of the executable, or `objdump -rdS <executable>` is needed to interpret this.</pre>			
/a/nojha-2018-06-21_00:18:52-rados-wip-24487-distro-basic-smithi/2686256			
Related issues:			
Duplicated by RADOS - Bug #25181: /mon/OSDMonitor.cc: 1821: FAILED assert(osd...		Duplicate	07/30/2018
Copied to RADOS - Backport #35071: mimic: FAILED assert(osdmap_manifest.pinne...		Resolved	

History

#1 - 06/27/2018 09:18 PM - Josh Durgin

- Priority changed from Normal to Urgent

#2 - 08/01/2018 01:06 PM - Sage Weil

- Status changed from New to Verified

/a/sage-2018-07-31_21:57:28-rados-wip-sage-testing-2018-07-31-1436-distro-basic-smithi/2844443
/a/sage-2018-07-30_13:46:50-rados-wip-sage3-testing-2018-07-28-1512-distro-basic-smithi/2838737

#3 - 08/01/2018 01:07 PM - Sage Weil

- Duplicated by Bug #25181: /mon/OSDMonitor.cc: 1821: FAILED assert(osdmap_manifest.pinned.empty()) added

#4 - 08/16/2018 05:58 PM - Sage Weil

/a/sage-2018-08-15_15:49:39-rados-wip-sage2-testing-2018-08-15-0731-distro-basic-smithi/2908178

#5 - 08/16/2018 06:09 PM - Joao Eduardo Luis

- Assignee set to Joao Eduardo Luis

#6 - 08/24/2018 12:54 PM - Joao Eduardo Luis

- Category set to Correctness/Safety

- Status changed from Verified to In Progress

#7 - 08/26/2018 12:50 PM - Joao Eduardo Luis

<https://github.com/ceph/ceph/pull/23742>

Currently missing: a reproducer. Reproducing may not be trivial because this requires a few conditions to be true to trigger the bug:

1. do_prune() must have been called, and the in-memory state **must** have been changed;
2. the transaction containing the update must be postponed because Paxos is busy;
3. an election must be called **before** the Paxos transaction is committed.

While triggering an election is trivial enough, ensuring we have the first two conditions is not, and would likely require instrumentation in the code to force the condition. I think we're better off relying on the suites, that seem to be reproducing this every now and then.

FWIW, the cause appears to be that we would not clear

#8 - 08/29/2018 05:25 PM - Sage Weil

- Status changed from In Progress to Pending Backport

- Backport set to mimic

#9 - 08/31/2018 11:07 AM - Nathan Cutler

- Copied to Backport #35071: mimic: FAILED assert(osdmap_manifest.pinned.empty()) in OSDMonitor::prune_init() added

#10 - 11/26/2018 11:49 PM - Nathan Cutler

- Status changed from Pending Backport to Resolved