

Ceph - Bug #19824

Reccurance of #18746(Jewel) in (Kraken)

05/02/2017 02:52 PM - Ross Martyn

Status:	Duplicate	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:			
Category:			
Target version:			
Source:		Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:		Pull request ID:	
Regression:	No	Crash signature (v1):	
Severity:	2 - major	Crash signature (v2):	
Reviewed:			
Description			
<p>Three months ago we had an issue where our monitors crashed following a run of a set of 'OpenStack refstack tests'. It appeared to be http://tracker.ceph.com/issues/18746</p> <p>Today we have had a very similar issue, twice, however this time during a user delete of a snapshot of a volume. Since the previous version of this issue we have upgraded to Kraken (11.2.0).</p> <p>We believe this to be very serious as the only way to rectify appeared to be rebooting all OpenStack infrastructure working with Ceph. (i.e all hypervisors & Cinder-volume instances) then restarting the monitors one by one. Far from ideal!</p> <pre>ceph version 11.2.0 (f223e27eeb35991352ebc1f67423d4ebc252adb7) 1: (()+0x6e77c2) [0x560058c2e7c2] 2: (()+0x10330) [0x7fc36c677330] 3: (gsignal()+0x37) [0x7fc36b322c37] 4: (abort()+0x148) [0x7fc36b326028] 5: (interval_set<snapid_t>::insert(snapid_t, snapid_t, snapid_t*, snapid_t*)+0x342) [0x560058b23a72] 6: (pg_pool_t::remove_unmanaged_snap(snapid_t)+0x4d) [0x560058b18a7d] 7: (OSDMonitor::prepare_pool_op(std::shared_ptr<MonOpRequest>)+0xde9) [0x5600588d1639] 8: (OSDMonitor::prepare_update(std::shared_ptr<MonOpRequest>)+0x32f) [0x5600588f7ccf] 9: (PaxosService::dispatch(std::shared_ptr<MonOpRequest>)+0xf02) [0x5600588a72e2] 10: (PaxosService::C_RetryMessage::_finish(int)+0x54) [0x5600588a8c74] 11: (C_MonOp::finish(int)+0x69) [0x5600588731f9] 12: (Context::complete(int)+0x9) [0x560058872369] 13: (void finish_contexts<Context>(CephContext*, std::list<Context*, std::allocator<Context*> >, int)+0x94) [0x560058879604] 14: (Paxos::finish_round()+0x10b) [0x56005889cc6b] 15: (Paxos::handle_last(std::shared_ptr<MonOpRequest>)+0xf86) [0x56005889e1b6] 16: (Paxos::dispatch(std::shared_ptr<MonOpRequest>)+0x2e4) [0x56005889eb64] 17: (Monitor::dispatch_op(std::shared_ptr<MonOpRequest>)+0xd19) [0x56005886cdb9] 18: (Monitor::_ms_dispatch(Message*)+0x6a1) [0x56005886d711] 19: (Monitor::ms_dispatch(Message*)+0x23) [0x56005888f2a3] 20: (DispatchQueue::entry()+0x793) [0x560058be9f63] 21: (DispatchQueue::DispatchThread::entry()+0xd) [0x560058aad8ad] 22: (()+0x8184) [0x7fc36c66f184] 23: (clone()+0x6d) [0x7fc36b3e9bed] NOTE: a copy of the executable, or `objdump -rDS <executable>` is needed to interpret this.</pre>			
Logs Attached.			
<ul style="list-style-type: none">Cinder Logs - Unfortunately not set to debug at the time.			

- Monitor Log - This is very long, but a lot happens in a few seconds!

Looking at the API log vs the Monitor log, The API call entry appears to be the request that killed the monitors (11:27:43).

There is no reference to this particular action in the Cinder Volume log.

For unfamiliar OpenStack users reference - Cinder API call -> Cinder Volume -> Ceph RBD image/snapshot 'Delete' call.

To make this slightly more complicated, during this fact finding mission we have found that pinned packages on our cinder controllers we pinning ceph-common and associated components to Jewel. Unsure if this miss match may be in part responsible. Looking to upgrade these now.

Related issues:

Duplicates RADOS - Bug #18746: monitors crashing ./include/interval_set.h: 35...

Resolved

01/30/2017

History

#1 - 06/17/2017 07:57 AM - Greg Farnum

- Duplicates Bug #18746: monitors crashing ./include/interval_set.h: 355: FAILED assert(0) (jewel+kraken) added

#2 - 06/17/2017 07:58 AM - Greg Farnum

- Status changed from New to Duplicate

Files

ceph-mon-1	525 KB	05/02/2017	Ross Martyn
Cinder-Vol	5.26 KB	05/02/2017	Ross Martyn