

Ceph - Bug #12637

bug of rados ls with keyvaluestore

08/06/2015 04:08 PM - Xinxin Shu

Status:	Resolved	Start date:	08/06/2015
Priority:	Normal	Due date:	
Assignee:	Xinxin Shu	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:	other	Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor		

Description

if leveldb is used as keyvaluestore backend , run 'rados -p rbd ls', then cluster crash , log shows as follow :

```
5> 2015-08-06 15:49:53.176711 7f38a4038700 10 osd.0 pg_epoch: 12 pg[0.0( empty local-les=12 n=0 ec=1 les/c 12/12 11/11/9) [0,1] r=0 lpr=11 crt=0'0 mlcod 0'0 active+clean] handle_message: 0x3b13d00
-----
4> 2015-08-06 15:49:53.176728 7f38a4038700 10 osd.0 pg_epoch: 12 pg[0.0( empty local-les=12 n=0 ec=1 les/c 12/12 11/11/9) [0,1] r=0 lpr=11 crt=0'0 mlcod 0'0 active+clean] do_pg_op osd_op(client.14108.0:1 [pgls start_epoch 0] 0.0 ack+read+known_if_redirected e12) v5
-----
3> 2015-08-06 15:49:53.176744 7f38a4038700 5
op tracker -- seq: 63, time: 2015-08-06 15:49:53.176744, event: started, op: osd_op(client.14108.0:1 [pgls start_epoch 0] 0.0 ack+read+known_if_redirected e12)
-2> 2015-08-06 15:49:53.176760 7f38a4038700 10 osd.0 pg_epoch: 12 pg[0.0( empty local-les=12 n=0 ec=1 les/c 12/12 11/11/9) [0,1] r=0 lpr=11 crt=0'0 mlcod 0'0 active+clean] pgls pg=0.0 count 1 024
-1> 2015-08-06 15:49:53.176781 7f38a4038700 10 collection_list_partial 0.0_head start:MIN/0 is_max:0
0> 2015-08-06 15:49:53.193495 7f38a4038700 -1 os/GenericObjectMap.cc: In function 'static std::string GenericObjectMap::header_key(const coll_t&, const ghobject_t&)' thread 7f38a4038700 time 2015-08-06 15:49:53.176820
os/GenericObjectMap.cc: 160: FAILED assert(oid.shard_id != shard_id_t::NO_SHARD)
```

ceph version 9.0.1-1988-gb63f2af (b63f2afadbc3b4436c936c5a9b62dcd416ea3afe)

```
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x8b) [0xb1443b]
2: (GenericObjectMap::header_key(coll_t const&, ghobject_t const&)+0x3e9) [0x9fe4c9]
3: (GenericObjectMap::list_objects(coll_t const&, ghobject_t, int, std::vector<ghobject_t, std::allocator<ghobject_t>&>, ghobject_t)+0x129) [0xa06fd9]
4: (KeyValueStore::collection_list_partial(coll_t, ghobject_t, int, int, snapid_t, std::vector<ghobject_t, std::allocator<ghobject_t>&>, ghobject_t)+0x1ff) [0x91b06f]
5: (PGBackend::objects_list_partial(hobject_t const&, int, int, snapid_t, std::vector<hobject_t, std::allocator<hobject_t>&>, hobject_t)+0x227) [0x8b05a7]
6: (ReplicatedPG::do_pg_op(std::tr1::shared_ptr<OpRequest>&)+0x102a) [0x85265a]
7: (ReplicatedPG::do_op(std::tr1::shared_ptr<OpRequest>&)+0x1eb) [0x877a2b]
8: (ReplicatedPG::do_request(std::tr1::shared_ptr<OpRequest>&+, ThreadPool::TPHandle&)+0x58c) [0x8155dc]
9: (OSD::dequeue_op(boost::intrusive_ptr<PG>&, std::tr1::shared_ptr<OpRequest>&+, ThreadPool::TPHandle&)+0x3bd) [0x6719cd]
10: (PGQueueable::RunVis::operator()(std::tr1::shared_ptr<OpRequest>&)+0x5d) [0x671c1d]
11: (OSD::ShardedOpWQ::_process(unsigned int, ceph::heartbeat_handle_d*)+0x8de) [0x69625e]
12: (ShardedThreadPool::shardedthreadpool_worker(unsigned int)+0x875) [0xb03e05]
13: (ShardedThreadPool::WorkThreadSharded::entry()+0x10) [0xb05f20]
14: (()+0x8182) [0x7f38d4344182]
15: (clone()+0x6d) [0x7f38d250447d]
```

NOTE: a copy of the executable, or `objdump -rdS <executable>` is needed to interpret this.

Associated revisions

Revision 62214832 - 08/07/2015 07:58 PM - xinxin shu

fix 'rados ls' issue if using keyvaluestore as objectstore backend

Fixes: #12637

Signed-off-by: xinxin shu <xinxin.shu@intel.com>

History

#1 - 08/06/2015 04:08 PM - Xinxin Shu

- Assignee set to Xinxin Shu

#2 - 08/06/2015 04:28 PM - Xinxin Shu

<https://github.com/ceph/ceph/pull/5499>

#3 - 09/07/2015 07:14 AM - Xinxin Shu

- Status changed from New to Resolved