

rbd - Bug #8564

osd cannot be restarted when leveldb is used as backend

06/09/2014 02:47 AM - Xinxin Shu

Status:	Resolved	Start date:	06/09/2014
Priority:	Urgent	Due date:	
Assignee:	Haomai Wang	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:	Community (dev)	Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature:	
Description			
<p>hi all , recently i enabled leveldb as filestore backend, after restarting my cluster , an osd is crashed, from the log , i get an error like this :</p> <pre>2014-06-09 16:22:12.250078 7f86a5610700 -1 os/KeyValueStore.cc: In function 'unsigned int KeyValueStore::_do_transaction(ObjectStore::Transaction&, KeyValueStore::BufferTransaction&, SequencerPosition&, ThreadPool::TPHandle*)' thread 7f86a5610700 time 2014-06-09 16:22:12.248673 os/KeyValueStore.cc: 1524: FAILED assert(0 == "unexpected error") ceph version 0.80-820-g5d606cd (5d606cd0d00698699c91a378a1bd9f71cc8a77c9) 1: (KeyValueStore::_do_transaction(ObjectStore::Transaction&, KeyValueStore::BufferTransaction&, SequencerPosition&, ThreadPool::TPHandle*)+0x750) [0x9e6f20] 2: (KeyValueStore::_do_transactions(std::list<ObjectStore::Transaction*, std::allocator<ObjectStore::Transaction*>> &, unsigned long, ThreadPool::TPHandle*)+0x8e) [0x9e8d1e] 3: (KeyValueStore::_do_op(KeyValueStore::OpSequencer*, ThreadPool::TPHandle&)+0x9a) [0x9e8e2a] 4: (ThreadPool::worker(ThreadPool::WorkThread*)+0x68a) [0xb62a3a] 5: (ThreadPool::WorkThread::entry()+0x10) [0xb63c90] 6: (()+0x7e9a) [0x7f86aed2de9a] 7: (clone()+0x6d) [0x7f86ad2d8ccd] NOTE: a copy of the executable, or `objdump -rds &lt;executable>` is needed to interpret this.</pre>			

History

#1 - 06/09/2014 04:29 AM - Haomai Wang

Hi xinxin,

Thanks for your report, you hint a known bug which will solved in (<https://github.com/ceph/ceph/pull/1649>) branch.

But I have cherry-picked the fix patches and push to PR. You can see commit message to know why(<https://github.com/yuyuyu101/ceph/commit/50c8fee8fda42f78ea563cab6229bdf0af3c8c99>). The PR is <https://github.com/ceph/ceph/pull/1941>

```
And the log shows the op dump: { "ops": [ { "op_num": 0,
"op_name": "remove",
"collection": "3.27_head",
"oid": "97a47827Vrbd_data.11756b8b4567.00000000000015f7VheadV3"}, { "op_num": 1,
"op_name": "mkcoll",
"collection": "3.27_TEMP"}, { "op_num": 2,
"op_name": "remove",
"collection": "3.27_TEMP",
"oid": "97a47827Vrbd_data.11756b8b4567.00000000000015f7VheadV3"}, { "op_num": 3,
"op_name": "touch",
"collection": "3.27_head",
"oid": "97a47827Vrbd_data.11756b8b4567.00000000000015f7VheadV3"}, { "op_num": 4,
"op_name": "omap_setheader",
"collection": "3.27_head",
"oid": "97a47827Vrbd_data.11756b8b4567.00000000000015f7VheadV3",
"header_length": "0"}, { "op_num": 5,
```

```
"op_name": "write",
"collection": "3.27_head",
"oid": "97a47827\rbdd_data.11756b8b4567.00000000000015f7\head\3",
"length": 4194304,
"offset": 0,
"bufferlist length": 4194304},    { "op_num": 6,
"op_name": "omap_setkeys",
"collection": "3.27_head",
"oid": "97a47827\rbdd_data.11756b8b4567.00000000000015f7\head\3",
"attr_lens": {}},    { "op_num": 7,
"op_name": "setattrs",
"collection": "3.27_head",
"oid": "97a47827\rbdd_data.11756b8b4567.00000000000015f7\head\3",
"attr_lens": { "_": 257,
"snapset": 31}},    { "op_num": 8,
"op_name": "omap_setkeys",
"collection": "meta",
"oid": "16ef7597\infos\head\1",
"attr_lens": { "3.27_epoch": 4,
"3.27_info": 684}}}}
/
```

#2 - 06/09/2014 04:30 AM - Haomai Wang

- Status changed from New to Need Review
- Assignee set to Haomai Wang
- Priority changed from Normal to Urgent
- Source changed from other to Community (dev)

#3 - 06/09/2014 10:48 PM - Xinxin Shu

hi haomai , is <https://github.com/ceph/ceph/pull/1941> PR the fix for this bug , while <https://github.com/ceph/ceph/pull/1941> PR is for the performance optimization

#4 - 06/09/2014 11:51 PM - Haomai Wang

Are you mean PR 1941 for bug fix and 1649 for performance purpose?
If so, yes

#5 - 06/16/2014 09:12 AM - Sage Weil

- Status changed from Need Review to Resolved

Files

osd.log	13.6 MB	06/09/2014	Xinxin Shu
---------	---------	------------	------------