

## Ceph - Bug #21769

assert(total\_data\_size % sinfo.get\_chunk\_size() == 0) with ec overwrite flag set

10/12/2017 07:00 AM - huang jun

<b>Status:</b>	Resolved	<b>Start date:</b>	10/12/2017
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	huang jun	<b>% Done:</b>	0%
<b>Category:</b>	OSD	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>		<b>Spent time:</b>	0.00 hour
<b>Source:</b>		<b>Reviewed:</b>	
<b>Tags:</b>		<b>Affected Versions:</b>	
<b>Backport:</b>	mimic, luminous	<b>ceph-qa-suite:</b>	
<b>Regression:</b>	No	<b>Pull request ID:</b>	
<b>Severity:</b>	2 - major		

### Description

scripts to reproduce:

```
#!/bin/sh
```

```
pool=rbd  
obj=object2
```

```
function set_size() {  
    local osd=$1  
    ./bin/ceph osd set noout  
    ./bin/init-ceph stop osd.$osd  
    dd if=/dev/urandom bs=10 count=1 of=CORRUPT  
    ./bin/ceph-objectstore-tool --data-path dev/osd$osd $obj set-bytes CORRUPT || return 1  
    ./bin/init-ceph start osd.$osd  
    ./bin/ceph tell osd.$osd injectargs '--debug_osd 20'  
    ./bin/ceph osd unset noout  
    ./bin/rados -p $pool get $obj /tmp/copy  
    diff /tmp/copy /tmp/9461760 || echo "Content not consistent"  
}
```

```
function create_cluster() {  
    ./src/vstart.sh -X -n --mon_num 1 --osd_num 4 --mds_num 0  
    ./bin/ceph osd erasure-code-profile set huangjun k=3 m=1  
    ./bin/ceph osd pool create $pool 1 1 erasure huangjun  
    ./bin/ceph osd pool set $pool allow_ec_overwrites true  
    dd if=/dev/zero of=/tmp/9461760 bs=1K count=947  
    truncate -s 9461760 /tmp/9461760  
    ./bin/rados -p $pool put $obj /tmp/9461760  
}
```

```
create_cluster
```

```
sed -i 's/bluestore fsck on mount = true/bluestore fsck on mount = false/g' ceph.conf  
./bin/ceph tell osd.* injectargs '--debug_osd 20'  
set_size 1  
set_size 0
```

```
0> 2017-10-12 12:20:30.127026 7fc9bf8d4700 -1 /usr/src/ceph/src/osd/ECUtil.cc: In function 'i  
nt ECUtil::decode(const ECUtil::stripe_info_t&, ceph::ErasureCodeInterfaceRef&, std  
::map<int, ceph::buffer::list>&, ceph::bufferlist*)' thread 7fc9bf8d4700 time 2017-10-12 12:20:30.  
118682  
/usr/src/ceph/src/osd/ECUtil.cc: 16: FAILED assert(total_data_size % sinfo.get_chunk_size() == 0)
```

```

ceph version 13.0.0-1208-g3d90ec5 (3d90ec5a4d6a29e4f2fea8fbc8abc532bc0801de) mimic (dev)
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x110) [0x7fc9dc5bcd0]
2: (ECUtil::decode(ECUtil::stripe_info_t const&, std::shared_ptr<ceph::ErasureCodeInterface>&, st
d::map<int, ceph::buffer::list, std::less<int>, std::allocator<std::pair<int const,
ceph::buffer::list> > >&, ceph::buffer::list*)+0x3f8) [0x7fc9dc35d608]
3: (CallClientContexts::finish(std::pair<RecoveryMessages*, ECBackend::read_result_t&&)+0x289) [
0x7fc9dc352e79]
4: (ECBackend::complete_read_op(ECBackend::ReadOp&, RecoveryMessages*)+0x7f) [0x7fc9dc329eaf]
5: (ECBackend::handle_sub_read_reply(pg_shard_t, ECSubReadReply&, RecoveryMessages*, ZTracer::Tra
ce const&)+0x1028) [0x7fc9dc331638]
6: (ECBackend::_handle_message(boost::intrusive_ptr<OpRequest>)+0x1af) [0x7fc9dc33b7bf]
7: (PGBackend::handle_message(boost::intrusive_ptr<OpRequest>)+0x50) [0x7fc9dc249d40]
8: (PrimaryLogPG::do_request(boost::intrusive_ptr<OpRequest>&, ThreadPool::TPHandle&)+0x5ae) [0x7
fc9dc1b8d1e]
9: (OSD::dequeue_op(boost::intrusive_ptr<PG>, boost::intrusive_ptr<OpRequest>, ThreadPool::TPHand
le&)+0x3f9) [0x7fc9dc056cb9]
10: (PGQueueable::RunVis::operator()(boost::intrusive_ptr<OpRequest> const&)+0x57) [0x7fc9dc2b4e1
7]
11: (OSD::ShardedOpWQ::_process(unsigned int, ceph::heartbeat_handle_d*)+0x1167) [0x7fc9dc0811e7]
12: (ShardedThreadPool::shardedthreadpool_worker(unsigned int)+0x839) [0x7fc9dc5c28c9]
13: (ShardedThreadPool::WorkThreadSharded::entry()+0x10) [0x7fc9dc5c4860]
14: (()+0x7dc5) [0x7fc9d93aldc5]
15: (clone()+0x6d) [0x7fc9d849573d]

```

#### Related issues:

Copied to Ceph - Backport #35959: mimic: assert(total_data_size % sinfo.get_c...	<b>Resolved</b>
Copied to Ceph - Backport #35960: luminous: assert(total_data_size % sinfo.ge...	<b>Resolved</b>

#### History

##### #1 - 10/12/2017 03:54 PM - huang jun

<https://github.com/ceph/ceph/pull/18266>

##### #2 - 10/17/2017 07:22 AM - Kefu Chai

- Status changed from New to Need Review
- Assignee set to huang jun

##### #3 - 09/05/2018 08:36 PM - David Zafman

- Backport set to mimic, luminous

It is possible that <https://github.com/ceph/ceph/pull/21611> resolves this in a better way.

##### #4 - 09/12/2018 04:47 PM - Sage Weil

- Status changed from Need Review to Pending Backport

<https://github.com/ceph/ceph/pull/21611>

##### #5 - 09/12/2018 07:26 PM - Nathan Cutler

- Copied to Backport #35959: mimic: assert(total\_data\_size % sinfo.get\_chunk\_size() == 0) with ec overwrite flag set added

##### #6 - 09/12/2018 07:26 PM - Nathan Cutler

- Copied to Backport #35960: luminous: assert(total\_data\_size % sinfo.get\_chunk\_size() == 0) with ec overwrite flag set added

##### #7 - 10/02/2018 08:08 PM - Yuri Weinstein

luminous backport PR <https://github.com/ceph/ceph/pull/24342> merged

**#8 - 10/08/2018 10:16 AM - Nathan Cutler**

*- Status changed from Pending Backport to Resolved*