RADOS - Bug #21573

[upgrade] buffer::list ABI broken in luminous release

09/27/2017 02:56 PM - Jason Dillaman

Status:	Resolved	**Start date:**	09/27/2017
Priority:	Urgent	**Due date:**	
Assignee:	Sage Weil	**% Done:**	0%
Category:		**Estimated time:**	0.00 hour
Target version:		**Spent time:**	0.00 hour
Source:		**Affected Versions:**	
Tags:		**ceph-qa-suite:**	
Backport:	luminous	**Component(RADOS):**	librados
Regression:	No	**Pull request ID:**	
Severity:	3 - minor	**Crash signature:**	
Reviewed:			

Description

A client application that was compiled against a pre-Luminous librados C++ API and therefore utilizing bufferlist will now be ABI incompatible w/ the Luminous version of librados. Specifically, the addition of "_mempool" within buffer::list causes a seg fault since pre-Luminous clients would not have initialized that field.

```
#0  mempool::pool_t::adjust_count (this=0x8737e8a5c980, items=items@entry=1, bytes=bytes@entry=400
8)
    at /home/jdillaman/ceph/src/common/mempool.cc:85
#1  0x00007ffff75dea47 in ceph::buffer::raw::reassign_to_mempool (this=<optimized out>, this=<opti
mized out>, pool=1879025072)
    at /home/jdillaman/ceph/src/common/buffer.cc:206
#2  ceph::buffer::list::append (this=this@entry=0x7fff6fffa5b0, data=data@entry=0x7fffee427477 "ce
ph v027", len=len@entry=9)
    at /home/jdillaman/ceph/src/common/buffer.cc:1912
#3  0x00007fffee0ab632 in AsyncConnection::_process_connection (this=this@entry=0x7fff64009730)
    at /home/jdillaman/ceph/src/msg/async/AsyncConnection.cc:923
#4  0x00007fffee0b1e98 in AsyncConnection::process (this=0x7fff64009730) at /home/jdillaman/ceph/s
rc/msg/async/AsyncConnection.cc:838
#5  0x00007fffee0c4405 in EventCenter::process_events (this=this@entry=0xb09910, timeout_microseco
nds=<optimized out>,
    timeout_microseconds@entry=30000000, working_dur=working_dur@entry=0x7fff6fffb2e8) at /home/jd
illaman/ceph/src/msg/async/Event.cc:436
#6  0x00007fffee0c8b68 in NetworkStack::<lambda()>::operator()(void) const (__closure=0xb369c8)
    at /home/jdillaman/ceph/src/msg/async/Stack.cc:53
#7  0x00007ffff703376f in ?? () from /lib64/libstdc++.so.6
#8  0x00007ffff730773a in start_thread () from /lib64/libpthread.so.0
#9  0x00007ffff6aa2e7f in clone () from /lib64/libc.so.6
```

Related issues:

Duplicated by Ceph - Bug #21352: librados version should bump major number wh...	**Duplicate**	**09/11/2017**
Copied to RADOS - Backport #21899: luminous: [upgrade] buffer::list ABI broke...	**Resolved**	

History

#1 - 09/27/2017 02:57 PM - Jason Dillaman

- Backport set to luminous

#2 - 09/27/2017 03:02 PM - Jason Dillaman

- Subject changed from [upgrade] bufferlist ABI broken in luminous release to [upgrade] buffer::list ABI broken in luminous release

- Description updated

#3 - 09/27/2017 03:02 PM - Jason Dillaman

- Description updated

#4 - 10/11/2017 12:12 AM - Jason Dillaman

- Priority changed from High to Urgent

#5 - 10/17/2017 08:41 PM - Yuri Weinstein

- Assignee set to Kefu Chai

@Kefu can you pls take a look?

#6 - 10/19/2017 09:14 AM - Kefu Chai

this would be a little bit tricky:

```
  class CEPH_BUFFER_API list {
    # ..
    int _mempool = -1;
    # ..
  public:
    // cons/des
    list() : _len(0), _memcopy_count(0), last_p(this) {}
    // cppcheck-suppress noExplicitConstructor
    list(unsigned prealloc) : _len(0), _memcopy_count(0), last_p(this) {
      reserve(prealloc);
    }
    // ...
  };
```

so the constructors are inlined, and are not defined in librados. that's why _mempool is not initialized by the old clients, which are still using their own copies of the ctors.

and we have **no** way to change the behavior of the inlined constructors.

so i'd suggest bump up the so version of librados. will send a mail to ceph-devel asking for more opinions.

#7 - 10/19/2017 09:30 AM - Kefu Chai

- Duplicated by Bug #21352: librados version should bump major number when ABI changes added

#8 - 10/20/2017 05:30 AM - Kefu Chai

- Status changed from New to Need Review

- Assignee changed from Kefu Chai to Sage Weil

- Component(RADOS) librados added

https://github.com/ceph/ceph/pull/18408

#9 - 10/23/2017 03:53 PM - Sage Weil

- Status changed from Need Review to Pending Backport

#10 - 10/24/2017 03:52 AM - Kefu Chai

- Copied to Backport #21899: luminous: [upgrade] buffer::list ABI broken in luminous release added

#11 - 10/24/2017 03:53 AM - Kefu Chai

- Status changed from Pending Backport to Resolved