

rgw - Bug #35905

multisite: update index segfault on shutdown/realm reload

09/10/2018 10:48 AM - Tianshan Qu

Status:	Resolved	Start date:	09/09/2018
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:	mimic luminous	ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature:	

Description

RGWAsyncFetchRemoteObj and RGWAsyncRemoveObj use zone_trace pointer from RGWBucketSyncSingleEntryCR. shutdown/realm reload will lead to CR Manager stop, but async rados op may still ongoing.

```
1: ((()+0x230722) [0x7f3a4d943722]
2: ((()+0xf100) [0x7f3a4c243100]
3: (gsignal()+0x37) [0x7f3a40f2b5f7]
4: (abort()+0x148) [0x7f3a40f2cce8]
5: (_gnu_cxx::_verbose_terminate_handler()+0x165) [0x7f3a4182fac5]
6: ((()+0x5ea36) [0x7f3a4182da36]
7: ((()+0x5ea63) [0x7f3a4182da63]
8: (__cxa_rethrow()+0x49) [0x7f3a4182dcd9]
9: (std::_Rb_tree<std::string, std::string, std::_Identity<std::string>;, std::less<std::string>;, std::allocator<std::string>; >::_M_copy(std::_Rb_tree_node<std::string>; const*, std::_Rb_tree_node<std::string>;*)+0x14d) [0x7f3a4d9aca7d]
10: (std::_Rb_tree<std::string, std::string, std::_Identity<std::string>;, std::less<std::string>;, std::allocator<std::string>; >::operator=(std::_Rb_tree<std::string, std::string, std::_Identity<std::string>;, std::less<std::string>;, std::allocator<std::string>; > const&)+0x47) [0x7f3a4d9acb27]
11: (RGWRados::cls_obj_prepare_op(RGWRados::BucketShard&, RGWModifyOp, std::string&, rgw_obj&, unsigned short, rgw_cls_meta_last_size*, std::set<std::string, std::less<std::string>;, std::allocator<std::string>; >, bool, unsigned char)+0xb8) [0x7f3a4da76ab8]
12: ((()+0x363ddb) [0x7f3a4da76ddb]
13: (RGWRados::Bucket::UpdateIndex::guard_reshard(RGWRados::BucketShard**, std::function<int (RGWRados::BucketShard*)>)+0x7d) [0x7f3a4dab474d]
14: (RGWRados::Bucket::UpdateIndex::prepare(RGWModifyOp, std::string const*, unsigned char*)+0x10d) [0x7f3a4dab4dbd]
15: (RGWRados::Object::Write::_do_write_meta(unsigned long, unsigned long, std::map<std::string, ceph::buffer::list, std::less<std::string>;, std::allocator<std::pair<std::string const, ceph::buffer::list>; >>&, bool, bool, void*)+0x761) [0x7f3a4daef2e1]
16: (RGWRados::Object::Write::write_meta(unsigned long, unsigned long, std::map<std::string, ceph::buffer::list, std::less<std::string>;, std::allocator<std::pair<std::string const, ceph::buffer::list>; >>&)+0x3b6) [0x7f3a4daf1916]
17: (RGWPutObjProcessor_Atomic::do_complete(unsigned long, std::string const&, std::chrono::time_point<ceph::time_detail::real_clock, std::chrono::duration<unsigned long, std::ratio<11, 10000000001>>>, std::chrono::time_point<ceph::time_detail::real_clock, std::chrono::duration<unsigned long, std::ratio<11, 10000000001>>>, std::map<std::string, ceph::buffer::list, std::less<std::string>;, std::allocator<std::pair<std::string const, ceph::buffer::list>; >>&, std::chrono::time_point<ceph::time_detail::real_clock, std::chrono::duration<unsigned long, std::ratio<11, 10000000001>>>, char const, char const*, std::string const*, std::set<std::string, std::less<std::string>;, std::allocator<std::string>; >, bool)+0x5cf) [0x7f3a4daf200f]
18: (RGWPutObjProcessor::complete(unsigned long, std::string const&, std::chrono::time_point<ceph::time_detail::real_clock, std::chrono::duration<unsigned long, std::ratio<11, 10000000001>>>
```

```

1&gt; > >
, std::chrono::time_point<ceph::time_detail::real_clock, std::chrono::duration<unsigned long
, std::ratio<11, 10000000001&gt; > >, std::map<std::string, ceph::buffer::list, std::less<
;std::string&gt;, std::allocator<std::pair<std::string const, ceph::buffer::list&gt; > >&, s
td::chrono::time_point<ceph::time_detail::real_clock, std::chrono::duration<unsigned long, s
td::ratio<11, 10000000001&gt; > >, char const*, char const*, std::string const*, std::set<st
d::string, std::less<std::string&gt;, std::allocator<std::string&gt; >
, bool)+0x4b) [0x7f3a4da5fefb]
19: (RGWRados::fetch_remote_obj(RGWObjectCtx&, rgw_user const&, std::string const&, std::string c
onst&, bool, req_info
, std::string const&, rgw_obj&, rgw_obj&, RGWBucketInfo&, RGWBucketInfo&, std::chrono::time_point&
lt;ceph::time_detail::real_clock, std::chrono::duration<unsigned long, std::ratio<11, 100000
00001&gt; > >
, std::chrono::time_point<ceph::time_detail::real_clock, std::chrono::duration<unsigned long
, std::ratio<11, 10000000001&gt; > >
, std::chrono::time_point<ceph::time_detail::real_clock, std::chrono::duration<unsigned long
, std::ratio<11, 10000000001&gt; > > const*, std::chrono::time_point<ceph::time_detail::real
_clock, std::chrono::duration<unsigned long, std::ratio<11, 10000000001&gt; > > const*, bool
, char const*, char const*, RGWRados::AttrsMod, bool, std::map<std::string, ceph::buffer::lis
t&gt; > >&, RGWObjCategory, unsigned long, std::chrono::time_point<ceph::time_detail::real_cloc
k, std::chrono::duration<unsigned long, std::ratio<11, 10000000001&gt; > >, std::string*, st
d::string*, ceph::buffer::list*, void ()(long, void
), void*, std::set<std::string, std::less<std::string&gt;, std::allocator<std::string&gt;
>, bool)+0xce3) [0x7f3a4dad61b3]
20: (RGWAsyncFetchRemoteObj::_send_request()+0x425) [0x7f3a4d9f55a5]
21: (RGWAsyncRadosProcessor::handle_request(RGWAsyncRadosRequest)+0x22) [0x7f3a4d9f1eb2]
22: (RGWAsyncRadosProcessor::RGWWQ::_process(RGWAsyncRadosRequest*, ThreadPool::TPHandle&)+0xd) [
0x7f3a4d9f1f7d]
23: (ThreadPool::worker(ThreadPool::WorkThread*)+0xa9e) [0x7f3a438448de]
24: (ThreadPool::WorkThread::entry()+0x10) [0x7f3a43845830]
25: ((()+0x7dc5) [0x7f3a4c23bdc5]
26: (clone()+0x6d) [0x7f3a40fec21d]
NOTE: a copy of the executable, or `objdump -rdS &lt;executable&gt;` is needed to interpret this.

```

Related issues:

Copied to rgw - Backport #36137: luminous: multisite: update index segfault o...	Resolved
Copied to rgw - Backport #36138: mimic: multisite: update index segfault on s...	Resolved

History

#1 - 09/10/2018 01:06 PM - Casey Bodley

- Status changed from New to Need Review
- Backport set to mimic luminous

<https://github.com/ceph/ceph/pull/24007>

#2 - 09/13/2018 05:43 PM - Casey Bodley

- Status changed from Need Review to Testing

#3 - 09/19/2018 02:50 PM - Casey Bodley

- Status changed from Testing to Pending Backport

#4 - 09/24/2018 11:00 AM - Nathan Cutler

- Copied to Backport #36137: luminous: multisite: update index segfault on shutdown/realm reload added

#5 - 09/24/2018 11:00 AM - Nathan Cutler

- Copied to Backport #36138: mimic: multisite: update index segfault on shutdown/realm reload added

#6 - 10/17/2018 08:34 PM - Nathan Cutler

- Status changed from Pending Backport to Resolved