

rgw - Bug #20906

multisite: FAILED assert(prev_iter != pos_to_prev.end()) in
RGWMetaSyncShardCR::collect_children()

08/03/2017 07:42 PM - Casey Bodley

Status:	Resolved	Start date:	08/03/2017
Priority:	High	Due date:	
Assignee:	Casey Bodley	% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:		Reviewed:	
Tags:		Affected Versions:	
Backport:	luminous	ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor		

Description

```
2017-08-03 15:19:13.189356 7f05bcff9700 20 cr:s=0x7f053c004a30:op=0x7f053c0f3d90:21RGWReadMDLogEnt
riesCR: operate()
2017-08-03 15:19:13.189370 7f05bcff9700 20 cr:s=0x7f053c004a30:op=0x7f053c02f0c0:18RGWMetaSyncShar
dCR: operate()
2017-08-03 15:19:13.189378 7f05bcff9700 20 meta sync: incremental_sync:1665: shard_id=0 log_entry:
 1_1501787951.872747_56.1:bucket.instance:haixjc-1:b2f89ca6-a10a-425c-ad79-82ba886bd6fe.4109.1:201
7-08-03 15:19:11.872747
2017-08-03 15:19:13.189418 7f05bcff9700 20 cr:s=0x7f053c083170:op=0x7f053c0f3d90:24RGWMetaSyncSing
leEntryCR: operate()
2017-08-03 15:19:13.189423 7f05bcff9700 20 meta sync: skipping pending operation
2017-08-03 15:19:13.189440 7f05bcff9700 20 cr:s=0x7f053c004a30:op=0x7f053c02f0c0:18RGWMetaSyncShar
dCR: operate()
2017-08-03 15:19:13.189454 7f05bcff9700 20 meta sync: incremental_sync:1665: shard_id=0 log_entry:
 1_1501787952.051313_57.1:bucket.instance:haixjc-1:b2f89ca6-a10a-425c-ad79-82ba886bd6fe.4109.1:201
7-08-03 15:19:12.051313
2017-08-03 15:19:13.189457 7f05bcff9700 0 meta sync: ERROR: cannot start syncing 1_1501787952.051
313_57.1. Duplicate entry?
2017-08-03 15:19:13.189460 7f05bcff9700 20 meta sync: incremental_sync:1665: shard_id=0 log_entry:
 1_1501787952.111111_58.1:bucket:haixjc-1:2017-08-03 15:19:12.111111
2017-08-03 15:19:13.189480 7f05bcff9700 20 cr:s=0x7f053c083170:op=0x7f053c0f3d90:24RGWMetaSyncSing
leEntryCR: operate()
2017-08-03 15:19:13.189487 7f05bcff9700 20 run: stack=0x7f053c083170 is done
2017-08-03 15:19:13.189518 7f05bcff9700 20 cr:s=0x7f053c071a20:op=0x7f053c0a4450:24RGWMetaSyncSing
leEntryCR: operate()
2017-08-03 15:19:13.189521 7f05bcff9700 20 meta sync: skipping pending operation
2017-08-03 15:19:13.189529 7f05bcff9700 20 cr:s=0x7f053c004a30:op=0x7f053c02f0c0:18RGWMetaSyncShar
dCR: operate()
2017-08-03 15:19:13.189535 7f05bcff9700 20 meta sync: incremental_sync:1665: shard_id=0 log_entry:
 1_1501787952.263821_59.1:bucket:haixjc-1:2017-08-03 15:19:12.263821
2017-08-03 15:19:13.189539 7f05bcff9700 0 meta sync: ERROR: cannot start syncing 1_1501787952.263
821_59.1. Duplicate entry?
2017-08-03 15:19:13.189548 7f05bcff9700 20 meta sync: incremental_sync:1665: shard_id=0 log_entry:
 1_1501787952.473367_60.1:bucket.instance:haixjc-2:b2f89ca6-a10a-425c-ad79-82ba886bd6fe.4109.2:201
7-08-03 15:19:12.473367
2017-08-03 15:19:13.189572 7f05bcff9700 20 cr:s=0x7f053c071a20:op=0x7f053c0a4450:24RGWMetaSyncSing
leEntryCR: operate()
2017-08-03 15:19:13.189580 7f05bcff9700 20 run: stack=0x7f053c071a20 is done
2017-08-03 15:19:13.189585 7f05bcff9700 20 cr:s=0x7f053c0726e0:op=0x7f053c0f3d90:24RGWMetaSyncSing
leEntryCR: operate()
2017-08-03 15:19:13.189587 7f05bcff9700 20 meta sync: skipping pending operation
2017-08-03 15:19:13.189593 7f05bcff9700 20 cr:s=0x7f053c004a30:op=0x7f053c02f0c0:18RGWMetaSyncShar
dCR: operate()
```

```
2017-08-03 15:19:13.189598 7f05bcfff9700 20 meta sync: incremental_sync:1665: shard_id=0 log_entry:
1_1501787952.611686_61.1:bucket:haixjc-2:b2f89ca6-a10a-425c-ad79-82ba886bd6fe.4109.2:201
7-08-03 15:19:12.611686
2017-08-03 15:19:13.189602 7f05bcfff9700 0 meta sync: ERROR: cannot start syncing 1_1501787952.611
686_61.1. Duplicate entry?
2017-08-03 15:19:13.189604 7f05bcfff9700 20 meta sync: incremental_sync:1665: shard_id=0 log_entry:
1_1501787952.660282_62.1:bucket:haixjc-2:2017-08-03 15:19:12.660282
2017-08-03 15:19:13.189624 7f05bcfff9700 20 cr:s=0x7f053c0726e0:op=0x7f053c0f3d90:24RGWMetaSyncSing
leEntryCR: operate()
2017-08-03 15:19:13.189630 7f05bcfff9700 20 run: stack=0x7f053c0726e0 is done
2017-08-03 15:19:13.189636 7f05bcfff9700 20 cr:s=0x7f053c068790:op=0x7f053c0a45c0:24RGWMetaSyncSing
leEntryCR: operate()
2017-08-03 15:19:13.189637 7f05bcfff9700 20 meta sync: skipping pending operation
2017-08-03 15:19:13.189643 7f05bcfff9700 20 cr:s=0x7f053c004a30:op=0x7f053c02f0c0:18RGWMetaSyncShar
dCR: operate()
2017-08-03 15:19:13.189648 7f05bcfff9700 20 meta sync: incremental_sync:1665: shard_id=0 log_entry:
1_1501787952.787951_65.1:bucket:haixjc-2:2017-08-03 15:19:12.787951
2017-08-03 15:19:13.189651 7f05bcfff9700 0 meta sync: ERROR: cannot start syncing 1_1501787952.787
951_65.1. Duplicate entry?
2017-08-03 15:19:13.189655 7f05bcfff9700 4 meta sync: cr:s=0x7f053c004a30:op=0x7f053c02f0c0:18RGWM
etaSyncShardCR: adjusting marker pos=1_1501787951.872747_56.1
2017-08-03 15:19:13.196648 7f05bcfff9700 -1 /home/cbodley/ceph/src/rgw/rgw_sync.cc: In function 'vo
id RGWMetaSyncShardCR::collect_children()' thread 7f05bcfff9700 time 2017-08-03 15:19:13.189662
/home/cbodley/ceph/src/rgw/rgw_sync.cc: 1398: FAILED assert(prev_iter != pos_to_prev.end())
```

```
ceph version Development (no_version) luminous (rc)
```

```
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x137) [0x7f0630aebbb2]
2: (RGWMetaSyncShardCR::collect_children()+0x303) [0x5569581341bb]
3: (RGWMetaSyncShardCR::incremental_sync()+0x2862) [0x556958138dea]
4: (RGWMetaSyncShardCR::operate()+0x1ee) [0x556958133cd6]
5: (RGWCoroutinesStack::operate(RGWCoroutinesEnv*)+0x191) [0x556957e7af39]
6: (RGWCoroutinesManager::run(std::__cxx11::list<RGWCoroutinesStack*, std::allocator<RGWCoroutine
sStack*> >)+0x290) [0x556957e7cb04]
7: (RGWCoroutinesManager::run(RGWCoroutine*)+0xbc) [0x556957e7deb0]
8: (RGWRemoteMetaLog::run_sync()+0x175e) [0x55695812045c]
9: (RGWMetaSyncStatusManager::run()+0x1c) [0x556957f7470a]
10: (RGWMetaSyncProcessorThread::process()+0x1c) [0x556957f76f06]
11: (RGWRadosThread::Worker::entry()+0xe8) [0x556957f10dd2]
12: (Thread::entry_wrapper()+0xc1) [0x7f0630f372a3]
13: (Thread::_entry_func(void*)+0x18) [0x7f0630f371d8]
14: (()+0x773a) [0x7f0639b5173a]
15: (clone()+0x3f) [0x7f062cf83e0f]
NOTE: a copy of the executable, or `objdump -rds <executable>` is needed to interpret this.
```

Further debugging shows that RGWMetaSyncShardCR is looping over the same entries a second time. The marker_tracker detects some of those duplicates (see the "Duplicate entry?" messages), but not others.

Related issues:

Copied to rgw - Backport #21097: luminous: multisite: FAILED assert(prev_iter...

Resolved

History

#1 - 08/03/2017 09:46 PM - Casey Bodley

possibly related to cls changes in <https://github.com/ceph/ceph/pull/16667?>

#2 - 08/15/2017 03:32 PM - Casey Bodley

- Status changed from New to Need Review

- Backport set to luminous

<https://github.com/ceph/ceph/pull/17024>

#3 - 08/17/2017 05:55 PM - Orit Wasserman

- Status changed from *Need Review* to *Need Test*

#4 - 08/17/2017 05:55 PM - Orit Wasserman

- Assignee set to *Casey Bodley*

#5 - 08/21/2017 03:56 PM - Yuri Weinstein

Casey Bodley wrote:

<https://github.com/ceph/ceph/pull/17024>

merged

#6 - 08/21/2017 04:48 PM - Nathan Cutler

- Status changed from *Need Test* to *Pending Backport*

#7 - 08/24/2017 04:06 PM - Abhishek Lekshmanan

- Copied to Backport #21097: *luminous: multisite: FAILED assert(prev_iter != pos_to_prev.end()) in RGWMetaSyncShardCR::collect_children() added*

#8 - 10/02/2017 01:50 PM - Nathan Cutler

- Status changed from *Pending Backport* to *Resolved*