

rbd - Bug #45694

"TestMigration.StressLive" fails

05/25/2020 09:39 AM - Kefu Chai

Status: Pending Backport	% Done: 0%
Priority: Normal	Spent time: 0.00 hour
Assignee: Jason Dillaman	
Category:	
Target version:	
Source:	Reviewed:
Tags:	Affected Versions:
Backport: nautilus,octopus	ceph-qa-suite:
Regression: No	Pull request ID: 38494
Severity: 3 - minor	Crash signature:
Description	
<pre>[RUN] TestMigration.StressLive ... /home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure Value of: src_bl.contents_equal(dst_bl) Actual: false Expected: true after commit snap: user1, block 29360128~4194304 differs /home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure Value of: src_bl.contents_equal(dst_bl) Actual: false Expected: true after commit snap: user2, block 29360128~4194304 differs /home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure Value of: src_bl.contents_equal(dst_bl) Actual: false Expected: true after commit snap: user3, block 29360128~4194304 differs /home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure Value of: src_bl.contents_equal(dst_bl) Actual: false Expected: true [FAILED] TestMigration.StressLive (5373 ms) ... The following tests FAILED: 6 - run-rbd-unit-tests-N.sh (Failed)</pre>	
found in 'make check' runs performed by jenkins.	
see https://github.com/ceph/ceph/pull/35222	
Related issues:	
Copied to rbd - Backport #48564: octopus: "TestMigration.StressLive" fails	New
Copied to rbd - Backport #48565: nautilus: "TestMigration.StressLive" fails	Need More Info
History	

#1 - 05/25/2020 09:39 AM - Kefu Chai

hi Jason, mind taking a look?

#2 - 05/25/2020 09:41 AM - Kefu Chai

- Description updated

#3 - 06/01/2020 03:44 PM - Kefu Chai

spotted again <https://github.com/ceph/ceph/pull/35337>

#4 - 08/04/2020 11:28 AM - Sebastian Wagner

```
[ RUN      ] TestMigration.StressLive
resize: 83886080
write: A 14189046~7453310
discard: 20225079~2701729
write: A 30885423~5709049
discard: 46481821~6993441
write: A 77667211~6218869
discard: 52802612~6566589
write: A 48946668~7555490
discard: 7683394~7533955
snap_create: snap0
resize: 150454613
write: B 102177830~2310820
discard: 103664482~4432103
write: B 73141758~6662509
discard: 35848104~5951569
write: B 51378288~2250537
discard: 2300898~704489
write: B 108821095~2247739
discard: 19822048~5940012
snap_create: snap1
write: C 145427126~3260454
discard: 147960604~1821887
write: C 30230407~110720
discard: 128410804~5977216
write: C 17303492~7216763
discard: 53317568~1727031
write: C 14438287~7100246
discard: 24027691~433762
snap_create: snap2
resize: 175529756
write: D 37093319~6348350
discard: 156757704~7225194
write: D 73302210~110316
discard: 47241027~3943201
write: D 43576944~2130965
discard: 19886125~6189903
write: D 3585797~6929446
discard: 119696274~972285
snap_create: snap3
migration_prepare
migration_execute
write: a 33578431~1083005
discard: 123629250~6449592
write: a 83829801~7240762
discard: 33684588~5391721
write: a 6578422~1364859
discard: 148074120~5797466
write: a 97267290~3808617
discard: 82844977~7399050
snap_create: user0
write: b 11429345~2249154
discard: 43532859~5790903
write: b 60146552~2028604
discard: 121687496~1516139
write: b 161964577~6613694
discard: 3614455~891649
write: b 94269726~419703
discard: 139709526~280376
```

```
snap_create: user1
write: c 170223077~3945066
discard: 164748816~2078231
write: c 42676630~4659561
discard: 159373483~7526602
write: c 41549139~4493790
discard: 142770056~750901
write: c 76922502~4615008
discard: 94696886~5753386
snap_create: user2
resize: 295079925
write: d 262162336~33718
discard: 192801602~5073272
write: d 39680065~1819219
discard: 2011973~755268
write: d 76138121~5549953
discard: 274429112~1296967
write: d 97033459~1544304
discard: 131925684~5289513
snap_create: user3
resize: 424312527
write: 1 42431252~10607813
discard: 45083205~5303906
write: 2 84862505~10607813
discard: 87514458~5303906
write: 3 127293758~10607813
discard: 129945711~5303906
write: 4 169725010~10607813
discard: 172376963~5303906
write: 5 212156263~10607813
discard: 214808216~5303906
migration_commit
after commit snap: user0, block 4194304~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
after commit snap: user1, block 4194304~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
after commit snap: user2, block 4194304~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
after commit snap: user3, block 4194304~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
after commit snap: null, block 4194304~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
[ FAILED ] TestMigration.StressLive (3571 ms)
```

<https://jenkins.ceph.com/job/ceph-pull-requests/56940/consoleFull#-78274084e840cee4-f4a4-4183-81dd-42855615f2c1>

```
[ RUN ] TestMigration.StressLive
resize: 83886080
write: A 13702922~2678541
discard: 13881004~6810979
write: A 49110948~6936813
discard: 51541425~4051199
write: A 53674946~6887683
discard: 18777008~2019298
write: A 82464644~210883
discard: 21455060~7998643
snap_create: snap0
write: B 35588620~5280154
discard: 78181009~5705071
write: B 64823389~4046205
discard: 5955270~7625100
write: B 540442~3117380
discard: 9789278~4777867
write: B 65035742~6723938
discard: 20426860~4079211
snap_create: snap1
resize: 143789335
write: C 97082460~2197849
discard: 111384147~5234198
write: C 30194323~2016233
discard: 84353109~5426925
write: C 120263003~399772
discard: 61539489~3626061
write: C 12562828~4852372
discard: 111435969~7111804
snap_create: snap2
resize: 134728747
write: D 112147305~8271968
discard: 37278495~336064
write: D 63834187~4657881
discard: 36036624~4401410
write: D 93601043~3040810
discard: 12590142~6642855
write: D 123215103~1275013
discard: 109428887~2907217
snap_create: snap3
migration_prepare
migration_execute
write: a 40252532~4923451
discard: 7917439~898041
write: a 105055808~7097872
discard: 40619730~4176930
write: a 45615540~1367295
discard: 86400172~2510076
write: a 96021109~1186877
discard: 68325028~3419490
snap_create: user0
resize: 174192390
write: b 38801862~1573590
discard: 3536609~2433032
write: b 127323974~5381676
discard: 67428748~2821049
write: b 43670898~3800059
discard: 166513437~7678953
write: b 2668508~7132870
discard: 8030911~230154
snap_create: user1
resize: 200261933
write: c 95085658~4545835
discard: 25884323~2726298
write: c 103899017~5247526
discard: 28534915~1060653
write: c 44422529~2127245
discard: 173514294~3350112
write: c 53254696~2511766
discard: 160198378~1695047
snap_create: user2
write: d 134363623~7076723
```

```
discard: 83215959~3922945
write: d 91210835~2881295
discard: 132889060~2750197
write: d 12287288~5940353
discard: 131540440~5032112
write: d 78533008~3340051
discard: 185377102~1739883
snap_create: user3
resize: 194842727
write: 1 19484272~4871068
discard: 20702039~2435534
write: 2 38968545~4871068
discard: 40186312~2435534
write: 3 58452818~4871068
discard: 59670585~2435534
write: 4 77937090~4871068
discard: 79154857~2435534
write: 5 97421363~4871068
discard: 98639130~2435534
migration_commit
after commit snap: user0, block 46137344~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
after commit snap: user0, block 67108864~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
after commit snap: user1, block 67108864~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
after commit snap: user2, block 67108864~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
after commit snap: user3, block 67108864~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
after commit snap: null, block 67108864~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
[ FAILED ] TestMigration.StressLive (2776 ms)
[-----] 54 tests from TestMigration (31173 ms total)
```

<https://jenkins.ceph.com/job/ceph-pull-requests/57111/consoleFull#512427994e840cee4-f4a4-4183-81dd-42855615f2c1>

#6 - 09/01/2020 08:08 PM - Neha Ojha

<https://jenkins.ceph.com/job/ceph-pull-requests/58934/console>

#7 - 10/12/2020 07:01 AM - Kefu Chai

```
[ RUN      ] TestMigration.StressLive
resize: 83886080
write: A 71159214~6293317
discard: 57226936~4585181

...
migration_commit
after commit snap: user0, block 134217728~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
after commit snap: user1, block 134217728~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
after commit snap: user2, block 134217728~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
after commit snap: user3, block 134217728~4194304 differs
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:153: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
[ FAILED  ] TestMigration.StressLive (4511 ms)
```

#8 - 12/04/2020 02:58 PM - Mykola Golub

```
+ RBD_FEATURES=0
+ unittest_librbd
seed 31644

[ RUN      ] TestMigration.StressLive
```

resize: 83886080
write: A 51954309~3684362
discard: 22699479~7618843
write: A 71241603~6061678
discard: 70164068~4760356
write: A 28177674~3356414
discard: 17048965~4326973
write: A 6335607~6523038
discard: 44355057~6453884
snap_create: snap0
resize: 115982112
write: B 51053858~4806985
discard: 48395937~1395353
write: B 56095267~921955
discard: 74284574~3477251
write: B 30524995~816379
discard: 25545521~3613463
write: B 99934305~8009634
discard: 25837278~7474166
snap_create: snap1
resize: 143005349
write: C 79902140~2984001
discard: 70257175~8202441
write: C 37296226~4215079
discard: 24722242~6612165
write: C 11469997~2087806
discard: 31860485~4761527
write: C 67091729~3219041
discard: 122054907~3529069
snap_create: snap2
resize: 161041485
write: D 99806025~4182622
discard: 114149849~4334734
write: D 147436554~3343624
discard: 100312080~2461899
write: D 127475520~6857415
discard: 110793202~1815912
write: D 68691854~6694694
discard: 73617392~4503518
snap_create: snap3
resize: 134547717
migration_prepare
migration_execute
write: a 117425823~2260727
discard: 117499188~2679499
write: a 128894998~5652719
discard: 72680734~8240768
write: a 16198578~232322
discard: 122963909~5097547
write: a 57961873~7832254
discard: 11837000~1474376
snap_create: user0
write: b 78835443~4818000
discard: 19610031~8352463
write: b 132099161~1806778
discard: 487221~1209955
write: b 64400135~8206007
discard: 6325927~1221977
write: b 117672522~372981
discard: 40981376~198046
snap_create: user1
resize: 243196365
write: c 196441668~7000978
discard: 142068194~4398739
write: c 239680160~3516205
discard: 86956611~4193070
write: c 22301450~7223616
discard: 221230066~4105305
write: c 51159413~2586588
discard: 204316637~5951360
snap_create: user2
resize: 310224549
write: d 203388968~4061756
discard: 130170116~4476597
write: d 73269279~7148563

```

discard: 267722339~1816474
write: d 86516388~7250541
discard: 3709186~481196
write: d 259756202~6561099
discard: 271314580~6377282
snap_create: user3
resize: 153137706
write: 1 15313770~3828442
discard: 16270880~1914221
write: 2 30627541~3828442
discard: 31584651~1914221
write: 3 45941311~3828442
discard: 46898421~1914221
write: 4 61255082~3828442
discard: 62212192~1914221
write: 5 76568853~3828442
discard: 77525963~1914221
migration_commit
after commit snap: user0, block 79691776~4194304 differs
src block: 4f2d3f50ae6:
00000000 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
*
0012c390 00 00 00 00 00 00 00 00 00 00 00 00 00 43 43 |.....CC|
0012c3a0 43 43 43 43 43 43 43 43 43 43 43 43 43 43 43 |CCCCCCCCCCCC|
*
0030bdf0 43 43 43 43 43 43 43 43 43 43 43 00 00 00 |CCCCCCCCCCCC...|
0030be00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
*
003ffff0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
00400000
dst block: 4f2ff52e2a4:
00000000 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
*
000335b0 00 00 00 00 00 00 00 00 00 00 00 00 43 43 43 |.....CCCC|
000335c0 43 43 43 43 43 43 43 43 43 43 43 43 43 43 |CCCCCCCCCCCC|
*
0030bdf0 43 43 43 43 43 43 43 43 43 43 43 00 00 00 |CCCCCCCCCCCC...|
0030be00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
*
003ffff0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
00400000
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:150: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
[ FAILED ] TestMigration.StressLive (4019 ms)

```


#9 - 12/04/2020 03:39 PM - Jason Dillaman

That last one is interesting. The simplified output of actions that could have affected the object is below:

```
[ RUN      ] TestMigration.StressLive
snap_create: snap0
snap_create: snap1
write: C 79902140~2984001 <----- WRITE 0 through 30BDFC
snap_create: snap2
snap_create: snap3
migration_prepare
migration_execute
discard: 72680734~8240768 <---- DISCARD 0 through 12C39E
snap_create: user0 <---- MISMATCH
write: b 78835443~4818000
snap_create: user1
snap_create: user2
write: d 73269279~7148563
snap_create: user3
write: 5 76568853~3828442
migration_commit
after commit snap: user0, block 79691776~4194304 differs
src block: 4f2d3f50ae6:
00000000 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
*
0012c390 00 00 00 00 00 00 00 00 00 00 00 00 43 43 |.....CC|
0012c3a0 43 43 43 43 43 43 43 43 43 43 43 43 43 43 |CCCCCCCCCCCC|
*
0030bdf0 43 43 43 43 43 43 43 43 43 43 43 00 00 00 |CCCCCCCCCCCC...|
0030be00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
*
003ffff0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
00400000
dst block: 4f2ff52e2a4:
00000000 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
*
000335b0 00 00 00 00 00 00 00 00 00 00 00 00 43 43 43 43 |.....CCCC|
000335c0 43 43 43 43 43 43 43 43 43 43 43 43 43 43 43 |CCCCCCCCCCCC|
*
0030bdf0 43 43 43 43 43 43 43 43 43 43 43 00 00 00 |CCCCCCCCCCCC...|
0030be00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
*
003ffff0 00 00 00 00 00 00 00 00 00 00 00 00 00 00 |.....|
00400000
/home/jenkins-build/build/workspace/ceph-pull-requests/src/test/librbd/test_Migration.cc:150: Failure
Value of: src_bl.contents_equal(dst_bl)
  Actual: false
Expected: true
[ FAILED ] TestMigration.StressLive (4019 ms)
```

Nothing should have been able to miss that discard from extent 0 -> 12C39E. We know that snapshot "snap3" had the correct data post-execute so that implies that the discard or the deep-copy write somehow only applied a partial extent.

#10 - 12/08/2020 02:57 AM - Jason Dillaman

- Status changed from New to In Progress

Finally found the bug. There is a race between a migration completing and a potential copy-up. If the copy-up starts while a migration is still in-progress, it will kick off the object deep-copy state machine. However, if it migration completes and the destination image refreshes before the in-flight deep-copy state machine issues the writes, it will incorrectly drop the "assert_snapc" guards and overwrite data previously written.

#11 - 12/08/2020 03:44 PM - Jason Dillaman

- Backport set to nautilus,octopus

#12 - 12/08/2020 07:19 PM - Jason Dillaman

- Status changed from In Progress to Fix Under Review

- Pull request ID set to 38494

#13 - 12/11/2020 03:56 PM - Mykola Golub

- Status changed from Fix Under Review to Pending Backport

#14 - 12/11/2020 04:00 PM - Backport Bot

- Copied to Backport #48564: octopus: "TestMigration.StressLive" fails added

#15 - 12/11/2020 04:00 PM - Backport Bot

- Copied to Backport #48565: nautilus: "TestMigration.StressLive" fails added