

Ceph - Bug #9891

"Assertion: os/DBObjectMap.cc: 1214: FAILED assert(0)" in upgrade:firefly-x-giant-distro-basic-multi run

10/26/2014 08:42 AM - Yuri Weinstein

Status:	Resolved	% Done:	0%
Priority:	Urgent	Spent time:	0.00 hour
Assignee:	Samuel Just		
Category:			
Target version:			
Source:	Q/A	Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:	giant,firefly	Pull request ID:	
Regression:	No	Crash signature (v1):	
Severity:	3 - minor	Crash signature (v2):	
Reviewed:			

Description

Logs are in http://qa-proxy.ceph.com/teuthology/teuthology-2014-10-25_18:13:01-upgrade:firefly-x-giant-distro-basic-multi/571584/

(from scrape)

```
Assertion: os/DBObjectMap.cc: 1214: FAILED assert(0)
ceph version 0.86-104-gb05efdd (b05efddb77290b86eb5c150776c761ab84f66f37)
 1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x8b) [0xb81feb]
 2: (DBObjectMap::lookup_parent(std::tr1::shared_ptr<DBObjectMap::_Header>)+0x798) [0xa2da28]
 3: (DBObjectMap::_clear(std::tr1::shared_ptr<DBObjectMap::_Header>, std::tr1::shared_ptr<KeyValue
DB::TransactionImpl>)+0xf4) [0xa30e74]
 4: (DBObjectMap::clear(ghobject_t const&, SequencerPosition const*)+0x202) [0xa35962]
 5: (FileStore::_lfn_unlink(coll_t, ghobject_t const&, SequencerPosition const&, bool)+0x16b) [0x8f
f01b]
 6: (FileStore::_remove(coll_t, ghobject_t const&, SequencerPosition const&)+0x8b) [0x8ff4ab]
 7: (FileStore::_do_transaction(ObjectStore::Transaction&, unsigned long, int, ThreadPool::TPHandl
e*)+0x2cc1) [0x90ff01]
 8: (FileStore::_do_transactions(std::list<ObjectStore::Transaction*, std::allocator<ObjectStore::
Transaction*> >&, unsigned long, ThreadPool::TPHandle*)+0x64) [0x913624]
 9: (FileStore::_do_op(FileStore::OpSequencer*, ThreadPool::TPHandle&)+0x18f) [0x9137cf]
10: (ThreadPool::worker(ThreadPool::WorkThread*)+0xa46) [0xb73176]
11: (ThreadPool::WorkThread::entry()+0x10) [0xb74220]
12: (()+0x8182) [0x7ffbf6ff8182]
13: (clone()+0x6d) [0x7ffbf5563fbd]
```

```
archive_path: /var/lib/teuthworker/archive/teuthology-2014-10-25_18:13:01-upgrade:firefly-x-giant-
distro-basic-multi/571584
```

```
branch: giant
```

```
description: upgrade:firefly-x/parallel/{0-cluster/start.yaml 1-firefly-install/firefly.yaml
 2-workload/{ec-rados-default.yaml rados_api.yaml rados_loadgenbig.yaml test_rbd_api.yaml
  test_rbd_python.yaml} 3-upgrade-sequence/upgrade-all.yaml 4-final-workload/{rados-snaps-few-obje
cts.yaml
  rados_loadgenmix.yaml rados_mon_thrash.yaml rbd_cls.yaml rbd_import_export.yaml
  rgw_swift.yaml} distros/ubuntu_14.04.yaml}
```

```
email: ceph-qa@ceph.com
```

```
job_id: '571584'
```

```
kernel: &id001
```

```
kdb: true
```

```
  sha1: distro
last_in_suite: false
machine_type: plana,burnupi,mira
name: teuthology-2014-10-25_18:13:01-upgrade:firefly-x-giant-distro-basic-multi
nuke-on-error: true
os_type: ubuntu
os_version: '14.04'
overrides:
  admin_socket:
    branch: giant
  ceph:
    conf:
      mon:
        debug mon: 20
        debug ms: 1
        debug paxos: 20
        mon warn on legacy crush tunables: false
      osd:
        debug filestore: 20
        debug journal: 20
        debug ms: 1
        debug osd: 20
    log-whitelist:
      - slow request
      - scrub mismatch
      - ScrubResult
    sha1: b05efddb77290b86eb5c150776c761ab84f66f37
  ceph-deploy:
    branch:
      dev: giant
    conf:
      client:
        log file: /var/log/ceph/ceph-$name.$pid.log
      mon:
        debug mon: 1
        debug ms: 20
        debug paxos: 20
        osd default pool size: 2
  install:
    ceph:
      sha1: b05efddb77290b86eb5c150776c761ab84f66f37
  s3tests:
    branch: giant
  workunit:
    sha1: b05efddb77290b86eb5c150776c761ab84f66f37
owner: scheduled_teuthology@teuthology
priority: 99
roles:
- - mon.a
  - mds.a
  - osd.0
  - osd.1
- - mon.b
  - mon.c
  - osd.2
  - osd.3
- - client.0
  - client.1
suite: upgrade:firefly-x
suite_branch: giant
suite_path: /var/lib/teuthworker/src/ceph-qa-suite_giant
targets:
  ubuntu@burnupi55.front.sepia.ceph.com: ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQCoNkQL+cTE/GCwShmnI
odptdNXomyURFXp3UtW2a5/5qqq/4ic1ujJ5KAxPSS+uJCP5u+RVpoX607WIT0RC90jPo4yQRU2CXk0AneE0rS2JK0DBz5KJlp
KXwjQ5hkU71uaXMkUDJE9YE1faZ3C92hnlyIixjSmWfNEdxcGWUCFRc1/wXNbJsp5wyTIy+RHqwxuy+IRp3Oa1qvrBm2u0wEx5
qCFUvlj4M1F+Oi7c531BWXRMlh9VbDa3scs2qc3Cu5foglCMe4B6FxoJ9J8ZPfdX/lI3qHlgRdnTzZB3fCCN1dc9ZOreAbL/8e
V+fHRMxG6uuCKp2b9M+scxh4puZ2J
```

```

ubuntu@plana29.front.sepia.ceph.com: ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQDgHl5M5raEaWWKwH7thuC
+Glhc6U1VLHpvIMDGMez1VOELXUB197ZayurxpWurUteDEaqevUqjUgi08lEL0dHW4G9ulst8B3FqHrAvIvCmhKwGslznrorZx
2Bq+/8mreE4ocMwft/R8sFzBsM3npgdgSqCAdDBSgI78S92WSlHGAqUz1A0iJoGdwiRjTKhCrI/tuIYyXUU9z/2vIR9bJTp0fw
t2S3z4LnpKdsMU8BGpJwm8CPrVj/odRDZ/epgEFBySb4OUq68QDXlg8RHnL1D+zFVsoi/uU5o9rW55EAo7KnPDcdNCQSSzmaur
wRaZxVSqhomF0kiU9oLlwlD+s7N
ubuntu@plana38.front.sepia.ceph.com: ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQCNqJmM2k0mJ+z0UPorFQE
Rz0n6Dm6kgc51D9b/Mxebr4tNb2JIhI0L23thzhnsWW/q9UzIyVFZriTFmjdrvHjmfqByoihjUiX2QnHF0jDskMuywSUM+hvhw
U+P7s+BSC5HZEWf/oQJ/w8WuLXOfsrumJu2cQc6ow3Zj+UzeGw00wab5knKwJDG70avfrs3wQjcm1RQWSqv+oVG4/CXQnC+TmA
MR1/2G42XBNjjf5v18M47M5zZKW+czP5r40dC5LmNU8lyvmvikxIaoxQS5n9SPSutuAF1PM6UnSy+Rq2RZNVl/yRVWYyUVOUBw
wFN8EarZ9icstCv2P1rzsMMeNYv
tasks:
- internal.lock_machines:
  - 3
  - plana,burnupi,mira
- internal.save_config: null
- internal.check_lock: null
- internal.connect: null
- internal.push_inventory: null
- internal.serialize_remote_roles: null
- internal.check_conflict: null
- internal.check_ceph_data: null
- internal.vm_setup: null
- kernel: *id001
- internal.base: null
- internal.archive: null
- internal.coredump: null
- internal.sudo: null
- internal.syslog: null
- internal.timer: null
- chef: null
- clock.check: null
- install:
  branch: firefly
- print: '**** done installing firefly'
- ceph:
  fs: xfs
- print: '**** done ceph'
- parallel:
  - workload
  - upgrade-sequence
- print: '**** done parallel'
- rados:
  clients:
  - client.1
  objects: 50
  op_weights:
    delete: 50
    read: 100
    rollback: 50
    snap_create: 50
    snap_remove: 50
    write: 100
  ops: 4000
- workunit:
  clients:
    client.1:
      - rados/load-gen-mix.sh
- mon_thrash:
  revive_delay: 20
  thrash_delay: 1
- workunit:
  clients:
    client.1:
      - rados/test.sh
- workunit:
  clients:
    client.1:

```

```

    - cls/test_cls_rbd.sh
- workunit:
  clients:
    client.1:
      - rbd/import_export.sh
  env:
    RBD_CREATE_ARGS: --new-format
- rgw:
  - client.1
- swift:
  client.1:
    rgw_server: client.1
teuthology_branch: master
tube: multi
upgrade-sequence:
  sequential:
    - install.upgrade:
      mon.a: null
      mon.b: null
    - ceph.restart:
      - mon.a
      - mon.b
      - mon.c
      - mds.a
      - osd.0
      - osd.1
      - osd.2
      - osd.3
verbose: true
worker_log: /var/lib/teuthworker/archive/worker_logs/worker.multi.3115
workload:
  sequential:
    - rados:
      clients:
        - client.0
      ec_pool: true
      objects: 50
      op_weights:
        append: 100
        copy_from: 50
        delete: 50
        read: 100
        rmatr: 25
        rollback: 50
        setattr: 25
        snap_create: 50
        snap_remove: 50
        write: 0
      ops: 4000
    - workunit:
      branch: firefly
      clients:
        client.0:
          - cls
    - workunit:
      branch: firefly
      clients:
        client.0:
          - rados/load-gen-big.sh
    - workunit:
      branch: firefly
      clients:
        client.0:
          - rbd/test_librbd.sh
    - workunit:
      branch: firefly
      clients:

```

```
client.0:
- rbd/test_librbd_python.sh
```

```
description: upgrade:firefly-x/parallel/{0-cluster/start.yaml 1-firefly-install/firefly.yaml
  2-workload/{ec-rados-default.yaml rados_api.yaml rados_loadgenbig.yaml test_rbd_api.yaml
  test_rbd_python.yaml} 3-upgrade-sequence/upgrade-all.yaml 4-final-workload/{rados-snaps-few-obje
cts.yaml
  rados_loadgenmix.yaml rados_mon_thrash.yaml rbd_cls.yaml rbd_import_export.yaml
  rgw_swift.yaml} distros/ubuntu_14.04.yaml}
duration: 2967.701498031616
failure_reason: ''wait_until_healthy''reached maximum tries (150) after waiting for
  900 seconds'
flavor: basic
owner: scheduled_teuthology@teuthology
status: fail
success: false
```

History

#1 - 10/27/2014 02:44 PM - Samuel Just

```
2014-10-25 18:56:23.243456 7fefdcc3e700 20 filestore dbobjectmap: seq is 485
2014-10-25 18:56:23.243559 7fefdd43f700 20 filestore dbobjectmap: seq is 486
2014-10-25 18:56:23.243602 7fefdd43f700 20 filestore dbobjectmap: seq is 487
2014-10-25 18:56:23.243628 7fefdd43f700 20 filestore set_header: setting seq 324
2014-10-25 18:56:23.243642 7fefdd43f700 20 filestore set_map_header: setting 485 oid
fb2248dc/burnupi5544081-40/45//3/18446744073709551615/2 parent seq 324
2014-10-25 18:56:23.243655 7fefdd43f700 20 filestore set_map_header: setting 486 oid fb2248dc/burnupi5544081-40/45//3/484/2 parent seq 324
...
2014-10-25 18:56:30.801289 7ffbf8345900 0 ceph version 0.86-104-gb05efdd (b05efddb77290b86eb5c150776c761ab84f66f37), process ceph-osd,
pid 18817
...
2014-10-25 18:56:31.104055 7ffbf8345900 20 filestore (init)dbobjectmap: seq is 485
...
2014-10-25 18:56:40.629704 7ffbe8e4d700 20 filestore(/var/lib/ceph/osd/ceph-2) objectmap clone
2014-10-25 18:56:40.629834 7ffbe8e4d700 20 filestore dbobjectmap: seq is 487
2014-10-25 18:56:40.629878 7ffbe8e4d700 20 filestore dbobjectmap: seq is 488
2014-10-25 18:56:40.629907 7ffbe8e4d700 20 filestore set_header: setting seq 453
2014-10-25 18:56:40.629926 7ffbe8e4d700 20 filestore set_map_header: setting 486 oid
ac9ea7a1/burnupi5544081-15/head//3/18446744073709551615/1 parent seq 453
2014-10-25 18:56:40.629959 7ffbe8e4d700 20 filestore set_map_header: setting 487 oid
ac9ea7a1/burnupi5544081-15/64//3/18446744073709551615/1 parent seq 453
...
2014-10-25 18:56:41.920041 7ffbe8e4d700 20 filestore remove_map_header: removing 486 oid fb2248dc/burnupi5544081-40/45//3/484/2
2014-10-25 18:56:41.920056 7ffbe8e4d700 20 filestore clear_header: clearing seq 486
...
2014-10-25 18:56:41.933829 7ffbe8e4d700 20 filestore remove_map_header: removing 489 oid ac9ea7a1/burnupi5544081-15/head//3/360/1
2014-10-25 18:56:41.933842 7ffbe964e700 10 filestore oid: 1113f7/pglog_3.as2/0//-1 not skipping op, *spos 7669.0.4
2014-10-25 18:56:41.933844 7ffbe8e4d700 20 filestore clear_header: clearing seq 489
2014-10-25 18:56:41.933847 7ffbe964e700 10 filestore > header.spos 0.0.0
2014-10-25 18:56:41.933890 7ffbe964e700 15 filestore(/var/lib/ceph/osd/ceph-2) _omap_setkeys meta/1113f7/pglog_3.as2/0//-1
2014-10-25 18:56:41.933932 7ffbe964e700 10 filestore oid: 1113f7/pglog_3.as2/0//-1 not skipping op, *spos 7669.0.5
2014-10-25 18:56:41.933937 7ffbe964e700 10 filestore > header.spos 0.0.0
2014-10-25 18:56:41.934078 7ffbe964e700 10 filestore(/var/lib/ceph/osd/ceph-2) -ERANGE, len is 802
2014-10-25 18:56:41.934107 7ffbe964e700 10 filestore(/var/lib/ceph/osd/ceph-2) -ERANGE, got 401
2014-10-25 18:56:41.934108 7ffbe8e4d700 20 filestore lookup_parent: parent 486 for seq 489
```

So, the transaction updating dbobjectmap seq didn't apply for some reason. It would have been submitted before the transaction actually writing out the nodes using it, so that's definitely odd since burnupi5544081-40/45//3/484/2 was apparently written out with seq 486. This must be either a filesystem thing or a leveldb thing. It's not clear to me how it could have been a filesystem thing though.

#2 - 10/29/2014 10:00 AM - Sage Weil

- Priority changed from Urgent to High

does not appear to be a ceph issue.. either bad disk or leveldb corruption or something. lowering priority.

#3 - 12/02/2014 01:43 PM - Samuel Just

- Status changed from New to Rejected

#4 - 01/12/2015 07:51 AM - Loic Dachary

http://pulpito.ceph.com/loic-2015-01-10_11:54:25-rados-giant-backports-testing-basic-vps/694324/

```
-6> 2015-01-11 09:18:46.785891 7f6d14685700 1 -- 10.214.140.199:6803/19206 <== osd.5 10.214.138.167:0/629
3 12 ==== osd_ping(ping e459 stamp 2015-01-11 09:18:46.774003) v2 ==== 47+0+0 (3607071168 0 0) 0x5d42a00 con 0
x5abd000
-5> 2015-01-11 09:18:46.785913 7f6d14685700 1 -- 10.214.140.199:6803/19206 --> 10.214.138.167:0/6293 -- o
sd_ping(ping_reply e459 stamp 2015-01-11 09:18:46.774003) v2 -- ?+0 0x5a4ec00 con 0x5abd000
-4> 2015-01-11 09:18:46.785929 7f6d14685700 20 osd.2 459 share_map_peer 0x5a45000 already has epoch 459
-3> 2015-01-11 09:18:46.785822 7f6d12e82700 1 -- 10.214.140.199:6802/19206 <== osd.5 10.214.138.167:0/629
3 12 ==== osd_ping(ping e459 stamp 2015-01-11 09:18:46.774003) v2 ==== 47+0+0 (3607071168 0 0) 0x5dbfa00 con 0
x5abd160
-2> 2015-01-11 09:18:46.785998 7f6d12e82700 1 -- 10.214.140.199:6802/19206 --> 10.214.138.167:0/6293 -- o
sd_ping(ping_reply e459 stamp 2015-01-11 09:18:46.774003) v2 -- ?+0 0x5499200 con 0x5abd160
-1> 2015-01-11 09:18:46.786011 7f6d12e82700 20 osd.2 459 share_map_peer 0x5a45000 already has epoch 459
0> 2015-01-11 09:18:46.804341 7f6d1c7c3700 -1 os/DBObjectMap.cc: In function 'DBObjectMap::Header DBObjec
tMap::lookup_parent(DBObjectMap::Header)' thread 7f6d1c7c3700 time 2015-01-11 09:18:46.775993
os/DBObjectMap.cc: 1214: FAILED assert(0)
```

```
ceph version 0.87-87-g0ea20e6 (0ea20e6c51208d6710f469454ab3f964bfa7c9d2)
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x7f) [0xaa04bf]
2: (DBObjectMap::lookup_parent(std::tr1::shared_ptr<DBObjectMap::_Header>)+0x925) [0x9790e5]
3: (DBObjectMap::_clear(std::tr1::shared_ptr<DBObjectMap::_Header>, std::tr1::shared_ptr<KeyValueDB::Transact
ionImpl>)+0xe5) [0x9792c5]
4: (DBObjectMap::clear(ghobject_t const&, SequencerPosition const&)+0x1a5) [0x97d325]
5: (FileStore::lfn_unlink(coll_t, ghobject_t const&, SequencerPosition const&, bool)+0x250) [0x8ab400]
6: (FileStore::_remove(coll_t, ghobject_t const&, SequencerPosition const&)+0x172) [0x8ab872]
7: (FileStore::_do_transaction(ObjectStore::Transaction&, unsigned long, int, ThreadPool::TPHandle*)+0x3bcc)
[0x8bbc5c]
8: (FileStore::_do_transactions(std::list<ObjectStore::Transaction*, std::allocator<ObjectStore::Transaction*
> >&, unsigned long, ThreadPool::TPHandle*)+0x64) [0x8be064]
9: (FileStore::_do_op(FileStore::OpSequencer*, ThreadPool::TPHandle&)+0x297) [0x8be317]
10: (ThreadPool::worker(ThreadPool::WorkThread*)+0x48e) [0xa900ee]
11: (ThreadPool::WorkThread::entry()+0x10) [0xa92e80]
12: (()+0x7e9a) [0x7f6d27eabe9a]
13: (clone()+0x6d) [0x7f6d268762ed]
NOTE: a copy of the executable, or `objdump -rdS <executable>` is needed to interpret this.
```

#5 - 02/16/2015 05:04 PM - Yuri Weinstein

- Status changed from Rejected to New

See in release giant v0.87.1 RC

Run: http://pulpito.front.sepia.ceph.com/teuthology-2015-02-12_17:29:32-rados-giant-distro-basic-multi/

Job: [754710]

Logs: http://qa-proxy.ceph.com/teuthology/teuthology-2015-02-12_17:29:32-rados-giant-distro-basic-multi/754710/

```
Assertion: os/DBObjectMap.cc: 1214: FAILED assert(0)
ceph version 0.87-159-gccb0914 (ccb0914f76da23acdd7374233cd1939ab80ef3c8)
 1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x8b) [0xb82c4b]
 2: (DBObjectMap::lookup_parent(std::tr1::shared_ptr<DBObjectMap::_Header>)+0x798) [0xa2e6d8]
 3: (DBObjectMap::_clear(std::tr1::shared_ptr<DBObjectMap::_Header>, std::tr1::shared_ptr<KeyValueDB::TransactionImpl>)+0xf4) [0xa31b24]
 4: (DBObjectMap::clear(ghobject_t const&, SequencerPosition const*)+0x202) [0xa36612]
 5: (FileStore::lfn_unlink(coll_t, ghobject_t const&, SequencerPosition const&, bool)+0x16b) [0x8ff55b]
 6: (FileStore::_remove(coll_t, ghobject_t const&, SequencerPosition const&)+0x8b) [0x8ff9eb]
 7: (FileStore::_do_transaction(ObjectStore::Transaction&, unsigned long, int, ThreadPool::TPHandle*)+0x2cc1) [0x910441]
 8: (FileStore::_do_transactions(std::list<ObjectStore::Transaction*, std::allocator<ObjectStore::Transaction* >, unsigned long, ThreadPool::TPHandle*)+0x64) [0x913b64]
 9: (FileStore::_do_op(FileStore::OpSequencer*, ThreadPool::TPHandle&)+0x18f) [0x913d0f]
10: (ThreadPool::worker(ThreadPool::WorkThread*)+0xa46) [0xb73dd6]
11: (ThreadPool::WorkThread::entry()+0x10) [0xb74e80]
12: (()+0x8182) [0x7f5c252de182]
13: (clone()+0x6d) [0x7f5c23848efd]
```

#6 - 02/16/2015 11:21 PM - Yuri Weinstein

- Backport set to giant

#7 - 02/20/2015 07:02 PM - Yuri Weinstein

- Priority changed from High to Urgent

#8 - 02/20/2015 07:54 PM - Yuri Weinstein

Verified that this test:

rados/thrash-erasure-code/{clusters/fixed-2.yaml fs/xf.s.yaml msgr-failures/few.yaml thrashers/default.yaml workloads/ec-rados-plugin=jerasure-k=2-m=1.yaml}

passed in run http://pulpito.front.sepia.ceph.com/teuthology-2015-02-17_14:43:58-rados-giant-distro-basic-multi/

Job/logs - http://qa-proxy.ceph.com/teuthology/teuthology-2015-02-17_14:43:58-rados-giant-distro-basic-multi/763604/

#9 - 02/20/2015 09:32 PM - Samuel Just

- Status changed from New to 7
- Assignee set to Samuel Just
- Backport changed from giant to hammer,giant,firefly,dumpling

2015-02-13 03:40:47.208535 7f0e9b5ec700 20 filestore dbobjectmap: seq is 271
2015-02-13 03:40:47.208564 7f0e9adeb700 20 filestore dbobjectmap: seq is 271
2015-02-13 03:40:47.208579 7f0e9b5ec700 20 filestore dbobjectmap: seq is 272

Note that there are two threads participating there! DBObjectMap::sync() doesn't take the header_lock!

#10 - 02/20/2015 09:32 PM - Samuel Just

This can be ignored for the present giant release though.

#11 - 03/04/2015 03:12 PM - Samuel Just

- Status changed from 7 to Pending Backport

#12 - 03/11/2015 09:18 AM - Loïc Dachary

- firefly backport <https://github.com/ceph/ceph/pull/3948>

#13 - 03/19/2015 09:32 AM - Loïc Dachary

- giant backport <https://github.com/ceph/ceph/pull/4094>

#14 - 03/26/2015 12:23 PM - Loïc Dachary

- Backport changed from hammer,giant,firefly,dumpling to hammer,giant,firefly

dumpling is end of life

#15 - 03/26/2015 12:27 PM - Loïc Dachary

- Backport changed from hammer,giant,firefly to giant,firefly

already in hammer

#16 - 03/26/2015 12:27 PM - Loïc Dachary

- Status changed from Pending Backport to Resolved