

fs - Bug #36349

mds: src/mds/MDCache.cc: 1637: FAILED ceph_assert(follows >= realm->get_newest_seq())

10/08/2018 11:06 PM - Patrick Donnelly

Status:	Can't reproduce	Start date:	10/08/2018
Priority:	Immediate	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:	v14.0.0	Affected Versions:	
Source:	Q/A	ceph-qa-suite:	
Tags:		Component(FS):	MDS
Backport:	mimic,luminous	Labels (FS):	crash
Regression:	No	Pull request ID:	
Severity:	3 - minor		
Reviewed:			

Description

```

2018-10-06T17:13:20.928 INFO:tasks.ceph.mds.l-s.smithi099.stderr:/build/ceph-14.0.0-3907-g276c86e/
src/mds/MDCache.cc: In function 'void MDCache::journal_cow_dentry(MutationImpl*, EMetaBlob*, CDent
ry*, snapid_t, CInode**, CDentry::linkage_t*)' thread 7fb68a8be700 time 2018-10-06 17:13:20.042739
2018-10-06T17:13:20.928 INFO:tasks.ceph.mds.l-s.smithi099.stderr:/build/ceph-14.0.0-3907-g276c86e/
src/mds/MDCache.cc: 1637: FAILED ceph_assert(follows >= realm->get_newest_seq())
2018-10-06T17:13:20.928 INFO:tasks.ceph.mds.l-s.smithi099.stderr: ceph version 14.0.0-3907-g276c86
e (276c86e4890a25e5e74b3f30dbb94987ede03b5a) nautilus (dev)
2018-10-06T17:13:20.928 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 1: (ceph::__ceph_assert_fail(char
 const*, char const*, int, char const*)+0x158) [0x7fb692749f4f]
2018-10-06T17:13:20.928 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 2: (ceph::__ceph_assertf_fail(char
 const*, char const*, int, char const*, char const*, ...) +0) [0x7fb69274a12c]
2018-10-06T17:13:20.928 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 3: (MDCache::journal_cow_dentry(
MutationImpl*, EMetaBlob*, CDentry*, snapid_t, CInode**, CDentry::linkage_t*)+0xee1) [0x5b33a1]
2018-10-06T17:13:20.928 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 4: (MDCache::journal_dirty_inode
(MutationImpl*, EMetaBlob*, CInode*, snapid_t)+0xc0) [0x5b3470]
2018-10-06T17:13:20.928 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 5: (MDCache::predirty_journal_pa
rents(boost::intrusive_ptr<MutationImpl>, EMetaBlob*, CInode*, CDir*, int, int, snapid_t)+0x631) [
0x5b9271]
2018-10-06T17:13:20.928 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 6: (Server::_unlink_local(boost:
:intrusive_ptr<MDRequestImpl>&, CDentry*, CDentry*)+0x65d) [0x5597ad]
2018-10-06T17:13:20.929 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 7: (Server::handle_client_unlink
(boost::intrusive_ptr<MDRequestImpl>&)+0xb76) [0x55f906]
2018-10-06T17:13:20.929 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 8: (Server::dispatch_client_requ
est(boost::intrusive_ptr<MDRequestImpl>&)+0xc55) [0x579585]
2018-10-06T17:13:20.929 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 9: (MDCache::dispatch_request(b
ost::intrusive_ptr<MDRequestImpl>&)+0x9c) [0x6207cc]
2018-10-06T17:13:20.929 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 10: (MDSInternalContextBase::com
plete(int)+0x72) [0x787432]
2018-10-06T17:13:20.929 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 11: (MDSCacheObject::finish_wait
ing(unsigned long, int)+0x283) [0x7a0db3]
2018-10-06T17:13:20.929 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 12: (SimpleLock::finish_waiters(
unsigned long, int)+0xb8) [0x691eb8]
2018-10-06T17:13:20.929 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 13: (Locker::eval_gather(SimpleL
ock*, bool, bool*, std::vector<MDSInternalContextBase*, std::allocator<MDSInternalContextBase*> >
)+0x1121) [0x67e601]
2018-10-06T17:13:20.929 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 14: (Locker::handle_file_lock(Sc
atterLock*, boost::intrusive_ptr<MLock const> const&)+0x9b3) [0x68f9c3]
2018-10-06T17:13:20.929 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 15: (Locker::handle_lock(boost:
intrusive_ptr<MLock const> const&)+0xa9) [0x690959]
2018-10-06T17:13:20.929 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 16: (Locker::dispatch(boost::int
rusive_ptr<Message const> const&)+0xbc) [0x690d0c]

```

```
2018-10-06T17:13:20.929 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 17: (MDSRank::handle_deferrable_
message(boost::intrusive_ptr<Message const> const&)+0x367) [0x4ea067]
2018-10-06T17:13:20.930 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 18: (MDSRank::_dispatch(boost::i
ntrusive_ptr<Message const> const&, bool)+0x6fb) [0x4eca0b]
2018-10-06T17:13:20.930 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 19: (MDSRankDispatcher::ms_dispa
tch(boost::intrusive_ptr<Message const> const&)+0x15) [0x4ed1b5]
2018-10-06T17:13:20.930 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 20: (MDSDaemon::ms_dispatch2(boo
st::intrusive_ptr<Message> const&)+0xfc) [0x4d9edc]
2018-10-06T17:13:20.930 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 21: (DispatchQueue::entry()+0x16
69) [0x7fb692978999]
2018-10-06T17:13:20.930 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 22: (DispatchQueue::DispatchThre
ad::entry()+0xd) [0x7fb692a258bd]
2018-10-06T17:13:20.930 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 23: (()+0x76ba) [0x7fb69200f6ba]
2018-10-06T17:13:20.930 INFO:tasks.ceph.mds.l-s.smithi099.stderr: 24: (clone()+0x6d) [0x7fb6918384
1d]
2018-10-06T17:13:20.930 INFO:tasks.ceph.mds.l-s.smithi099.stderr:*** Caught signal (Aborted) **
```

From:
/ceph/teuthology-archive/pdonnell-2018-10-06_01:15:54-multimds-wip-pdonnell-testing-20181005.225845-distro-basic-smithi/310688
0/teuthology.log

Core:
/ceph/teuthology-archive/pdonnell-2018-10-06_01:15:54-multimds-wip-pdonnell-testing-20181005.225845-distro-basic-smithi/310688
0/remote/smithi099/coredump/1538846000.16983.core

Branch: <https://github.com/batrick/ceph/tree/wip-pdonnell-testing-20181005.225845>

Related issues:

Related to Messengers - Bug #36540: msg: messages are queued but not sent **New**

History

#1 - 10/09/2018 12:26 PM - Zheng Yan

Looks like that mds_table_request(snaptable server_ready) got lost. It's the first message that mds.0 sent to mds.3

```
2018-10-06 17:11:23.844 7fec0858a700 1 -- 172.21.15.99:6817/2278901200 _send_to--> mds 172.21.15.99:6818/1206
97649 -- mds_table_request(snaptable server_ready) v1 -- ?+0 0x1546d00
```

...

```
2018-10-06 17:11:23.844 7fec0858a700 1 -- 172.21.15.99:6817/2278901200 <== mds.3 172.21.15.99:6818/120697649
2 ==== discover(1 0x1.* ) v1 ==== 35+0+0 (3238743458 0 0) 0x1546d00 con 0x2253c00
```

...

```
2018-10-06 17:11:23.844 7fec0858a700 1 -- 172.21.15.99:6817/2278901200 --> 172.21.15.99:6818/120697649 -- dis
cover_reply(1 0x1) v2 -- 0x21c0580 con 0
```

Can't find corresponding message in mds.3's log. The first message received from mds.0 is discover reply

```
2018-10-06 17:11:23.844 7fb68a8be700 1 -- 172.21.15.99:6818/120697649 <== mds.0 172.21.15.99:6817/2278901200
1 ==== discover_reply(1 0x1) v2 ==== 910+0+0 (3908223777 0 0) 0x2e34a00 con 0x2ecb100
```

#2 - 10/19/2018 10:46 PM - Patrick Donnelly

- Related to Bug #36540: *msg: messages are queued but not sent added*

#3 - 01/17/2019 04:15 PM - Patrick Donnelly

- Status changed from New to *Can't reproduce*

Haven't seen this since. Closing as can't reproduce. Probably noise from messenger changes?