

fs - Bug #41935

ceph mdss keep on crashing

09/19/2019 02:21 PM - Kenneth Waegeman

Status: Duplicate	% Done: 0%
Priority: Normal	
Assignee:	
Category:	
Target version:	
Source: Community (user)	Affected Versions: v14.2.3
Tags:	ceph-qa-suite:
Backport:	Component(FS): MDS
Regression:	Labels (FS): crash, multimds
Severity: 1 - critical	Pull request ID:
Reviewed:	Crash signature:

Description

I updated ceph to 14.2.3 yesterday. everything was running fine, but today mds start crashing. I tried restarting all mds, but issue remains

```
[root@mds03 ~]# ceph -s
cluster:
id: 92bfcf0a-1d39-43b3-b60f-44f01b630e47
health: HEALTH_WARN
insufficient standby MDS daemons available

services:
  mon: 3 daemons, quorum mds01,mds02,mds03 (age 22h)
  mgr: mds02(active, since 25h), standbys: mds01, mds03
  mds: ceph_fs:2 {0=mds02=up:active,1=mds03=up:active(laggy or crashed)}
  osd: 535 osds: 535 up, 535 in

data:
  pools: 3 pools, 3328 pgs
  objects: 370.95M objects, 666 TiB
  usage: 1.0 PiB used, 2.2 PiB / 3.2 PiB avail
  pgs: 3317 active+clean
        9 active+clean+scrubbing+deep
        2 active+clean+scrubbing
```

This is the stacktraces I see in th mds logs:

Related issues:

Duplicates fs - Bug #41948: nautilus: mds: incomplete backport of #40444 (MDC...	Resolved
--	-----------------

History

#1 - 09/19/2019 02:23 PM - Kenneth Waegeman

```
2019-09-19 15:41:40.156 7f037e6e01c0 0 set uid:gid to 167:167 (ceph:ceph)
2019-09-19 15:41:40.156 7f037e6e01c0 0 ceph version 14.2.3 (0f776cf838a1ae3130b2b73dc26be9c95c6ccc39) nautilus
s (stable), process ceph-mds, pid 294068
2019-09-19 15:41:40.156 7f037e6e01c0 0 pidfile_write: ignore empty --pid-file
2019-09-19 15:41:40.196 7f036c2f0700 1 mds.mds03 Updating MDS map to version 553402 from mon.2
2019-09-19 15:41:40.386 7f036c2f0700 1 mds.mds03 Updating MDS map to version 553403 from mon.2
2019-09-19 15:41:40.386 7f036c2f0700 1 mds.mds03 Map has assigned me to become a standby
2019-09-19 15:41:40.386 7f036c2f0700 1 mds.mds03 Updating MDS map to version 553404 from mon.2
```

```

2019-09-19 15:41:40.386 7f036c2f0700 1 mds.1.553404 handle_mds_map i am now mds.1.553404
2019-09-19 15:41:40.386 7f036c2f0700 1 mds.1.553404 handle_mds_map state change up:boot --> up:replay
2019-09-19 15:41:40.386 7f036c2f0700 1 mds.1.553404 replay_start
2019-09-19 15:41:40.386 7f036c2f0700 1 mds.1.553404 recovery set is 0
2019-09-19 15:41:40.386 7f036c2f0700 1 mds.1.553404 waiting for osdmap 203605 (which blacklists prior instance)
2019-09-19 15:41:40.656 7f0365ae3700 0 mds.1.cache creating system inode with ino:0xl01
2019-09-19 15:41:40.656 7f0365ae3700 0 mds.1.cache creating system inode with ino:0xl
2019-09-19 15:41:46.166 7f0364ae1700 1 mds.1.553404 Finished replaying journal
2019-09-19 15:41:46.166 7f0364ae1700 1 mds.1.553404 making mds journal writeable
2019-09-19 15:41:46.276 7f036c2f0700 1 mds.mds03 Updating MDS map to version 553405 from mds.0
2019-09-19 15:41:46.276 7f036c2f0700 1 mds.1.553404 handle_mds_map i am now mds.1.553404
2019-09-19 15:41:46.276 7f036c2f0700 1 mds.1.553404 handle_mds_map state change up:replay --> up:resolve
2019-09-19 15:41:46.276 7f036c2f0700 1 mds.1.553404 resolve_start
2019-09-19 15:41:46.276 7f036c2f0700 1 mds.1.553404 reopen_log
2019-09-19 15:41:46.276 7f036c2f0700 1 mds.1.553404 recovery set is 0
2019-09-19 15:41:46.566 7f036c2f0700 1 mds.1.553404 resolve_done
2019-09-19 15:41:47.406 7f036c2f0700 1 mds.mds03 Updating MDS map to version 553406 from mon.2
2019-09-19 15:41:47.406 7f036c2f0700 1 mds.1.553404 handle_mds_map i am now mds.1.553404
2019-09-19 15:41:47.406 7f036c2f0700 1 mds.1.553404 handle_mds_map state change up:resolve --> up:reconnect
2019-09-19 15:41:47.406 7f036c2f0700 1 mds.1.553404 reconnect_start
2019-09-19 15:41:47.406 7f036c2f0700 1 mds.1.server reconnect_clients -- 20 sessions
2019-09-19 15:41:47.406 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37962144 10.141.1.6.17:0/1354686824 after 0
2019-09-19 15:41:47.406 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37962162 10.141.1.6.24:0/2772929523 after 0
2019-09-19 15:41:47.406 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37962153 10.141.1.6.22:0/7426576 after 0
2019-09-19 15:41:47.406 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37962180 10.141.1.6.23:0/2413712232 after 0
2019-09-19 15:41:47.406 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37962183 10.141.1.6.20:0/743246641 after 0
2019-09-19 15:41:47.406 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37962189 10.141.1.6.19:0/2209083713 after 0
2019-09-19 15:41:47.406 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37962150 10.141.1.6.15:0/2231346138 after 0
2019-09-19 15:41:47.406 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37962177 10.141.1.6.25:0/648554407 after 0
2019-09-19 15:41:47.416 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.36165377 v1:10.141.1.1.10.90:0/1063780365 after 0.01
2019-09-19 15:41:47.416 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37960382 10.141.1.6.11:0/1996585949 after 0.01
2019-09-19 15:41:47.416 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37961209 10.141.1.6.21:0/2395019382 after 0.01
2019-09-19 15:41:47.416 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37960403 10.141.1.6.12:0/704451057 after 0.01
2019-09-19 15:41:47.416 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37961197 10.141.1.6.13:0/3937945744 after 0.01
2019-09-19 15:41:47.416 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37961236 10.141.1.6.14:0/304725641 after 0.01
2019-09-19 15:41:47.416 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37960373 10.141.1.6.18:0/2156995594 after 0.01
2019-09-19 15:41:47.416 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37960400 10.141.1.6.16:0/2737464876 after 0.01
2019-09-19 15:41:47.416 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37865319 10.141.1.6.21:0/1591580297 after 0.01
2019-09-19 15:41:47.416 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37865333 10.141.1.6.20:0/3660809459 after 0.01
2019-09-19 15:41:47.416 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37865372 10.141.1.6.19:0/3498217528 after 0.01
2019-09-19 15:41:47.456 7f036c2f0700 0 log_channel(cluster) log [DBG] : reconnect by client.37865356 10.141.1.6.11:0/1948991158 after 0.0500002
2019-09-19 15:41:47.456 7f036c2f0700 1 mds.1.553404 reconnect_done
2019-09-19 15:41:48.476 7f036c2f0700 1 mds.mds03 Updating MDS map to version 553407 from mon.2
2019-09-19 15:41:48.476 7f036c2f0700 1 mds.1.553404 handle_mds_map i am now mds.1.553404
2019-09-19 15:41:48.476 7f036c2f0700 1 mds.1.553404 handle_mds_map state change up:reconnect --> up:rejoin
2019-09-19 15:41:48.476 7f036c2f0700 1 mds.1.553404 rejoin_start
2019-09-19 15:41:48.476 7f036c2f0700 1 mds.1.553404 rejoin_joint_start
2019-09-19 15:41:49.216 7f036c2f0700 1 mds.1.553404 rejoin_done
2019-09-19 15:41:49.526 7f036c2f0700 1 mds.mds03 Updating MDS map to version 553408 from mon.2
2019-09-19 15:41:49.526 7f036c2f0700 1 mds.1.553404 handle_mds_map i am now mds.1.553404
2019-09-19 15:41:49.526 7f036c2f0700 1 mds.1.553404 handle_mds_map state change up:rejoin --> up:active
2019-09-19 15:41:49.526 7f036c2f0700 1 mds.1.553404 recovery_done -- successful recovery!
2019-09-19 15:41:49.526 7f036c2f0700 1 mds.1.553404 active_start
2019-09-19 15:41:49.756 7f036c2f0700 1 mds.1.553404 cluster recovered.

```

```
2019-09-19 15:42:00.336 7f036c2f0700 -1 /home/jenkins-build/build/workspace/ceph-build/ARCH/x86_64/AVAILABLE_ARCH/x86_64/AVAILABLE_DIST/centos7/DIST/centos7/MACHINE_SIZE/huge/release/14.2.3/rpm/el7/BUILD/ceph-14.2.3/src/mds/MDCache.cc: In function 'CInode* MDCache::cow_inode(CInode*, snapid_t)' thread 7f036c2f0700 time 2019-09-19 15:42:00.333567
/home/jenkins-build/build/workspace/ceph-build/ARCH/x86_64/AVAILABLE_ARCH/x86_64/AVAILABLE_DIST/centos7/DIST/centos7/MACHINE_SIZE/huge/release/14.2.3/rpm/el7/BUILD/ceph-14.2.3/src/mds/MDCache.cc: 1498: FAILED ceph_assert (!lock->get_num_wrllocks ())
```

```
ceph version 14.2.3 (0f776cf838a1ae3130b2b73dc26be9c95c6ccc39) nautilus (stable)
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x14a) [0x7f0375773ac2]
2: (ceph::__ceph_assertf_fail(char const*, char const*, int, char const*, char const*, ...) +0) [0x7f0375773c90]
3: ((()+0x1ee48d) [0x55a7e4ccf48d]
4: (MDCache::journal_cow_dentry(MutationImpl*, EMetaBlob*, CDentry*, snapid_t, CInode**, CDentry::linkage_t*)+0x823) [0x55a7e4ccfc3]
5: (MDCache::journal_dirty_inode(MutationImpl*, EMetaBlob*, CInode*, snapid_t)+0xbc) [0x55a7e4cd042c]
6: (Locker::_do_cap_update(CInode*, Capability*, int, snapid_t, boost::intrusive_ptr<MClientCaps const> const&, boost::intrusive_ptr<MClientCaps> const&, bool*)+0xfb6) [0x55a7e4d957e6]
7: (Locker::handle_client_caps(boost::intrusive_ptr<MClientCaps const> const&)+0x2059) [0x55a7e4d9c8e9]
8: (Locker::dispatch(boost::intrusive_ptr<Message const> const&)+0xe7) [0x55a7e4daaf97]
9: (MDSRank::handle_deferrable_message(boost::intrusive_ptr<Message const> const&)+0x304) [0x55a7e4c089e4]
10: (MDSRank::_dispatch(boost::intrusive_ptr<Message const> const&, bool)+0x6eb) [0x55a7e4c0b1bb]
11: (MDSRankDispatcher::ms_dispatch(boost::intrusive_ptr<Message const> const&)+0x40) [0x55a7e4c0b8d0]
12: (MDSDaemon::ms_dispatch2(boost::intrusive_ptr<Message> const&)+0xfc) [0x55a7e4bf82ac]
13: (DispatchQueue::entry()+0x12a9) [0x7f0375968dd9]
14: (DispatchQueue::DispatchThread::entry()+0xd) [0x7f0375a183dd]
15: ((()+0x7dd5) [0x7f0373642dd5]
16: (clone()+0x6d) [0x7f03722f302d]
```

```
2019-09-19 15:42:00.336 7f036c2f0700 -1 *** Caught signal (Aborted) **
in thread 7f036c2f0700 thread_name:ms_dispatch
```

```
ceph version 14.2.3 (0f776cf838a1ae3130b2b73dc26be9c95c6ccc39) nautilus (stable)
1: ((()+0xf5d0) [0x7f037364a5d0]
2: (gsignal()+0x37) [0x7f037222b2c7]
3: (abort()+0x148) [0x7f037222c9b8]
4: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x199) [0x7f0375773b11]
5: (ceph::__ceph_assertf_fail(char const*, char const*, int, char const*, char const*, ...) +0) [0x7f0375773c90]
6: ((()+0x1ee48d) [0x55a7e4ccf48d]
7: (MDCache::journal_cow_dentry(MutationImpl*, EMetaBlob*, CDentry*, snapid_t, CInode**, CDentry::linkage_t*)+0x823) [0x55a7e4ccfc3]
8: (MDCache::journal_dirty_inode(MutationImpl*, EMetaBlob*, CInode*, snapid_t)+0xbc) [0x55a7e4cd042c]
9: (Locker::_do_cap_update(CInode*, Capability*, int, snapid_t, boost::intrusive_ptr<MClientCaps const> const&, boost::intrusive_ptr<MClientCaps> const&, bool*)+0xfb6) [0x55a7e4d957e6]
10: (Locker::handle_client_caps(boost::intrusive_ptr<MClientCaps const> const&)+0x2059) [0x55a7e4d9c8e9]
11: (Locker::dispatch(boost::intrusive_ptr<Message const> const&)+0xe7) [0x55a7e4daaf97]
12: (MDSRank::handle_deferrable_message(boost::intrusive_ptr<Message const> const&)+0x304) [0x55a7e4c089e4]
13: (MDSRank::_dispatch(boost::intrusive_ptr<Message const> const&, bool)+0x6eb) [0x55a7e4c0b1bb]
14: (MDSRankDispatcher::ms_dispatch(boost::intrusive_ptr<Message const> const&)+0x40) [0x55a7e4c0b8d0]
15: (MDSDaemon::ms_dispatch2(boost::intrusive_ptr<Message> const&)+0xfc) [0x55a7e4bf82ac]
16: (DispatchQueue::entry()+0x12a9) [0x7f0375968dd9]
17: (DispatchQueue::DispatchThread::entry()+0xd) [0x7f0375a183dd]
18: ((()+0x7dd5) [0x7f0373642dd5]
19: (clone()+0x6d) [0x7f03722f302d]
NOTE: a copy of the executable, or `objdump -rds <executable>` is needed to interpret this.
```

```
--- begin dump of recent events ---
```

```
....
-221> 2019-09-19 15:41:48.476 7f036c2f0700 1 mds.mds03 Updating MDS map to version 553407 from mon.2
-220> 2019-09-19 15:41:48.476 7f036c2f0700 1 mds.1.553404 handle_mds_map i am now mds.1.553404
-219> 2019-09-19 15:41:48.476 7f036c2f0700 1 mds.1.553404 handle_mds_map state change up:reconnect --> up:r
ejoin
-218> 2019-09-19 15:41:48.476 7f036c2f0700 1 mds.1.553404 rejoin_start
-217> 2019-09-19 15:41:48.476 7f036c2f0700 1 mds.1.553404 rejoin_joint_start
-216> 2019-09-19 15:41:48.476 7f036c2f0700 5 mds.mds03 handle_mds_map old map epoch 553407 <= 553407, disca
rding
-215> 2019-09-19 15:41:48.486 7f036faf7700 5 mds.beacon.mds03 received beacon reply up:rejoin seq 5 rtt 1.0
3
-214> 2019-09-19 15:41:48.896 7f03652e2700 5 mds.1.log _submit_thread 30520209192609~1084 : ESessions 20 op
ens cmav 115357868
-213> 2019-09-19 15:41:49.216 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37962144
-212> 2019-09-19 15:41:49.216 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37960373
```

```

-211> 2019-09-19 15:41:49.216 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37960400
-210> 2019-09-19 15:41:49.216 7f036c2f0700 1 mds.1.553404 rejoin_done
-209> 2019-09-19 15:41:49.336 7f036c2f0700 3 mds.1.553404 request_state up:active
-208> 2019-09-19 15:41:49.336 7f036c2f0700 5 mds.beacon.mds03 set_want_state: up:rejoin -> up:active
-207> 2019-09-19 15:41:49.336 7f036c2f0700 5 mds.beacon.mds03 Sending beacon up:active seq 6
-206> 2019-09-19 15:41:49.336 7f036c2f0700 10 monclient: _send_mon_message to mon.mds03 at v2:10.141.16.3:33
00/0
-205> 2019-09-19 15:41:49.336 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37962150
-204> 2019-09-19 15:41:49.336 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37962177
-203> 2019-09-19 15:41:49.336 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37962162
-202> 2019-09-19 15:41:49.336 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37962153
-201> 2019-09-19 15:41:49.336 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37962180
-200> 2019-09-19 15:41:49.336 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37962183
-199> 2019-09-19 15:41:49.336 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37962189
-198> 2019-09-19 15:41:49.336 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37960403
-197> 2019-09-19 15:41:49.336 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37960382
-196> 2019-09-19 15:41:49.336 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37961197
-195> 2019-09-19 15:41:49.336 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37961236
-194> 2019-09-19 15:41:49.336 7f036c2f0700 3 mds.1.server handle_client_session client_session(request_rene
wcaps seq 9) v1 from client.37961209
-193> 2019-09-19 15:41:49.526 7f036c2f0700 1 mds.mds03 Updating MDS map to version 553408 from mon.2
-192> 2019-09-19 15:41:49.526 7f036c2f0700 1 mds.1.553404 handle_mds_map i am now mds.1.553404
-191> 2019-09-19 15:41:49.526 7f036c2f0700 1 mds.1.553404 handle_mds_map state change up:rejoin --> up:acti
ve
-190> 2019-09-19 15:41:49.526 7f036c2f0700 1 mds.1.553404 recovery_done -- successful recovery!
-189> 2019-09-19 15:41:49.526 7f036c2f0700 1 mds.1.553404 active_start
-188> 2019-09-19 15:41:49.526 7f036c2f0700 4 mds.1.553404 set_osd_epoch_barrier: epoch=203604
-187> 2019-09-19 15:41:49.526 7f03652e2700 5 mds.1.log _submit_thread 30520209193713~25997 : EUpdate cap up
date [metablob 0x1000bd8ac7b, 12 dirs]
-186> 2019-09-19 15:41:49.526 7f03652e2700 5 mds.1.log _submit_thread 30520209219730~1083 : EUpdate cap upd
ate [metablob 0x20005af5c74, 1 dirs]
-185> 2019-09-19 15:41:49.526 7f03652e2700 5 mds.1.log _submit_thread 30520209220833~1037 : EUpdate check_i
node_max_size [metablob 0x20005af5c74, 1 dirs]
-184> 2019-09-19 15:41:49.526 7f036faf7700 5 mds.beacon.mds03 received beacon reply up:active seq 6 rtt 0.1
90001
-183> 2019-09-19 15:41:49.756 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
319:2375 lookup #0x1001172f39d/koi8-c 2019-09-19 15:41:39.482105 RETRY=1 caller_uid=0, caller_gid=0{0,}) v4
-182> 2019-09-19 15:41:49.756 7f036c2f0700 5 mds.1.server waiting for root
-181> 2019-09-19 15:41:49.756 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8866 lookup #0x100166004d4/Windows-MSVC-C.cmake 2019-09-19 15:41:39.479661 RETRY=1 caller_uid=0, caller_gi
d=0{0,}) v4
-180> 2019-09-19 15:41:49.756 7f036c2f0700 5 mds.1.server waiting for root
-179> 2019-09-19 15:41:49.756 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
372:5807 setattr size=5055 #0x20005a6eb66 2019-09-19 15:41:39.478859 RETRY=1 caller_uid=0, caller_gid=0{0,}) v
4
-178> 2019-09-19 15:41:49.756 7f036c2f0700 5 mds.1.server waiting for root
-177> 2019-09-19 15:41:49.756 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
372:5808 lookup #0x20005a6eb3a/schema.cpython-37.pyc 2019-09-19 15:41:39.589633 RETRY=1 caller_uid=0, caller_g
id=0{0,}) v4
-176> 2019-09-19 15:41:49.756 7f036c2f0700 5 mds.1.server waiting for root
-175> 2019-09-19 15:41:49.756 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
356:3604570 readdir #0x2000090d629 2019-09-19 15:41:39.523863 RETRY=1 caller_uid=0, caller_gid=0{0,}) v4
-174> 2019-09-19 15:41:49.756 7f036c2f0700 5 mds.1.server waiting for root
-173> 2019-09-19 15:41:49.756 7f036c2f0700 1 mds.1.553404 cluster recovered.
-172> 2019-09-19 15:41:49.756 7f036c2f0700 4 mds.1.553404 set_osd_epoch_barrier: epoch=203605
-171> 2019-09-19 15:41:49.756 7f036c2f0700 5 mds.mds03 handle_mds_map old map epoch 553408 <= 553408, disca
rding
-170> 2019-09-19 15:41:50.196 7f036b2ee700 10 monclient: tick
-169> 2019-09-19 15:41:50.196 7f036b2ee700 10 monclient: _check_auth_rotating have uptodate secrets (they ex
pire after 2019-09-19 15:41:20.200249)
-168> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client log_queue is 20 last_log 20 sent 0 num 20 unsent 2
0 sending 20
-167> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.416276 mds.mds03 (md

```

```

s.1) 1 : cluster [DBG] reconnect by client.37962144 10.141.16.17:0/1354686824 after 0
-166> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.416356 mds.mds03 (md
s.1) 2 : cluster [DBG] reconnect by client.37962162 10.141.16.24:0/2772929523 after 0
-165> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.416403 mds.mds03 (md
s.1) 3 : cluster [DBG] reconnect by client.37962153 10.141.16.22:0/7426576 after 0
-164> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.416441 mds.mds03 (md
s.1) 4 : cluster [DBG] reconnect by client.37962180 10.141.16.23:0/2413712232 after 0
-163> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.416483 mds.mds03 (md
s.1) 5 : cluster [DBG] reconnect by client.37962183 10.141.16.20:0/743246641 after 0
-162> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.416524 mds.mds03 (md
s.1) 6 : cluster [DBG] reconnect by client.37962189 10.141.16.19:0/2209083713 after 0
-161> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.416562 mds.mds03 (md
s.1) 7 : cluster [DBG] reconnect by client.37962150 10.141.16.15:0/2231346138 after 0
-160> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.416625 mds.mds03 (md
s.1) 8 : cluster [DBG] reconnect by client.37962177 10.141.16.25:0/648554407 after 0
-159> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.416915 mds.mds03 (md
s.1) 9 : cluster [DBG] reconnect by client.36165377 v1:10.141.10.90:0/1063780365 after 0.01
-158> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.417130 mds.mds03 (md
s.1) 10 : cluster [DBG] reconnect by client.37960382 10.141.16.11:0/1996585949 after 0.01
-157> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.417177 mds.mds03 (md
s.1) 11 : cluster [DBG] reconnect by client.37961209 10.141.16.21:0/2395019382 after 0.01
-156> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.417212 mds.mds03 (md
s.1) 12 : cluster [DBG] reconnect by client.37960403 10.141.16.12:0/704451057 after 0.01
-155> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.417259 mds.mds03 (md
s.1) 13 : cluster [DBG] reconnect by client.37961197 10.141.16.13:0/3937945744 after 0.01
-154> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.417706 mds.mds03 (md
s.1) 14 : cluster [DBG] reconnect by client.37961236 10.141.16.14:0/304725641 after 0.01
-153> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.417815 mds.mds03 (md
s.1) 15 : cluster [DBG] reconnect by client.37960373 10.141.16.18:0/2156995594 after 0.01
-152> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.417935 mds.mds03 (md
s.1) 16 : cluster [DBG] reconnect by client.37960400 10.141.16.16:0/2737464876 after 0.01
-151> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.418627 mds.mds03 (md
s.1) 17 : cluster [DBG] reconnect by client.37865319 10.141.16.21:0/1591580297 after 0.01
-150> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.419295 mds.mds03 (md
s.1) 18 : cluster [DBG] reconnect by client.37865333 10.141.16.20:0/3660809459 after 0.01
-149> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.419728 mds.mds03 (md
s.1) 19 : cluster [DBG] reconnect by client.37865372 10.141.16.19:0/3498217528 after 0.01
-148> 2019-09-19 15:41:50.196 7f036b2ee700 10 log_client will send 2019-09-19 15:41:47.466316 mds.mds03 (md
s.1) 20 : cluster [DBG] reconnect by client.37865356 10.141.16.11:0/1948991158 after 0.0500002
-147> 2019-09-19 15:41:50.196 7f036b2ee700 10 monclient: _send_mon_message to mon.mds03 at v2:10.141.16.3:33
00/0
-146> 2019-09-19 15:41:50.216 7f0369aeb700 2 mds.1.cache Memory usage: total 2071464, rss 1637780, heap 32
3836, baseline 323836, 477 / 613577 inodes have caps, 482 caps, 0.000785557 caps per inode
-145> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209221890~9062 : EUpdate scatter
_writebehind [metablob 0x1000bd8ac7b, 2 dirs]
-144> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209230972~4965 : EUpdate scatter
_writebehind [metablob 0x1000bf27a9b, 2 dirs]
-143> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209235957~4519 : EUpdate scatter
_writebehind [metablob 0x100171e181c, 2 dirs]
-142> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209240496~2809 : EUpdate scatter
_writebehind [metablob 0x101, 2 dirs]
-141> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209243325~2809 : EUpdate scatter
_writebehind [metablob 0x101, 2 dirs]
-140> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209246154~4071 : EUpdate scatter
_writebehind [metablob 0x1000bf26309, 2 dirs]
-139> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209250245~10579 : EUpdate scatte
r_writebehind [metablob 0x1000bf26309, 9 dirs]
-138> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209260844~2305 : EUpdate scatter
_writebehind [metablob 0x1000bf2745b.001*, 2 dirs]
-137> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209263169~1845 : EUpdate scatter
_writebehind [metablob 0x1000c233753, 2 dirs]
-136> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209265034~1359 : EUpdate scatter
_writebehind [metablob 0x20005f881ed, 2 dirs]
-135> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209266413~5669 : EUpdate scatter
_writebehind [metablob 0x10010dcb97f.101*, 5 dirs]
-134> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209272102~1356 : EUpdate scatter
_writebehind [metablob 0x20005f881ef, 2 dirs]
-133> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209273478~1357 : EUpdate scatter
_writebehind [metablob 0x20005f881ed, 2 dirs]
-132> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209274855~1357 : EUpdate scatter
_writebehind [metablob 0x20005f881ed, 2 dirs]
-131> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209276232~1353 : EUpdate scatter
_writebehind [metablob 0x20005f881ed, 2 dirs]
-130> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209277605~1818 : EUpdate scatter
_writebehind [metablob 0x20005af5c68, 2 dirs]

```

```

-129> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209279443~1811 : EUpdate scatter
_writebehind [metablob 0x20005af5c69, 2 dirs]
-128> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209281274~12149 : EUpdate scatter
r_writebehind [metablob 0x1000bd8ac7b, 3 dirs]
-127> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209293443~9062 : EUpdate scatter
_writebehind [metablob 0x1000bd8ac7b, 2 dirs]
-126> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209302525~8615 : EUpdate scatter
_writebehind [metablob 0x1000bf27a9b, 2 dirs]
-125> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209311160~7042 : EUpdate scatter
_writebehind [metablob 0x1000bf2745a, 8 dirs]
-124> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209318222~2305 : EUpdate scatter
_writebehind [metablob 0x1000bf2745b.001*, 2 dirs]
-123> 2019-09-19 15:41:50.216 7f03652e2700 5 mds.1.log _submit_thread 30520209320547~3305 : EUpdate scatter
_writebehind [metablob 0x1000c233753, 2 dirs]
-122> 2019-09-19 15:41:50.226 7f03652e2700 5 mds.1.log _submit_thread 30520209323872~1846 : EUpdate scatter
_writebehind [metablob 0x20005f881ed, 2 dirs]
-121> 2019-09-19 15:41:50.226 7f03652e2700 5 mds.1.log _submit_thread 30520209325738~4038 : EUpdate scatter
_writebehind [metablob 0x20005af5c63, 3 dirs]
-120> 2019-09-19 15:41:50.226 7f03652e2700 5 mds.1.log _submit_thread 30520209329796~2760 : EUpdate scatter
_writebehind [metablob 0x20005af5c68, 2 dirs]
-119> 2019-09-19 15:41:50.226 7f03652e2700 5 mds.1.log _submit_thread 30520209332576~4651 : EUpdate lock in
est accounted scatter stat update [metablob 0x1000bd8ac7b, 1 dirs]
-118> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client handle_log_ack log(last 20) v1
-117> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.416276 mds.mds03 (mds.1
) 1 : cluster [DBG] reconnect by client.37962144 10.141.16.17:0/1354686824 after 0
-116> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.416356 mds.mds03 (mds.1
) 2 : cluster [DBG] reconnect by client.37962162 10.141.16.24:0/2772929523 after 0
-115> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.416403 mds.mds03 (mds.1
) 3 : cluster [DBG] reconnect by client.37962153 10.141.16.22:0/7426576 after 0
-114> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.416441 mds.mds03 (mds.1
) 4 : cluster [DBG] reconnect by client.37962180 10.141.16.23:0/2413712232 after 0
-113> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.416483 mds.mds03 (mds.1
) 5 : cluster [DBG] reconnect by client.37962183 10.141.16.20:0/743246641 after 0
-112> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.416524 mds.mds03 (mds.1
) 6 : cluster [DBG] reconnect by client.37962189 10.141.16.19:0/2209083713 after 0
-111> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.416562 mds.mds03 (mds.1
) 7 : cluster [DBG] reconnect by client.37962150 10.141.16.15:0/2231346138 after 0
-110> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.416625 mds.mds03 (mds.1
) 8 : cluster [DBG] reconnect by client.37962177 10.141.16.25:0/648554407 after 0
-109> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.416915 mds.mds03 (mds.1
) 9 : cluster [DBG] reconnect by client.36165377 v1:10.141.10.90:0/1063780365 after 0.01
-108> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.417130 mds.mds03 (mds.1
) 10 : cluster [DBG] reconnect by client.37960382 10.141.16.11:0/1996585949 after 0.01
-107> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.417177 mds.mds03 (mds.1
) 11 : cluster [DBG] reconnect by client.37961209 10.141.16.21:0/2395019382 after 0.01
-106> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.417212 mds.mds03 (mds.1
) 12 : cluster [DBG] reconnect by client.37960403 10.141.16.12:0/704451057 after 0.01
-105> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.417259 mds.mds03 (mds.1
) 13 : cluster [DBG] reconnect by client.37961197 10.141.16.13:0/3937945744 after 0.01
-104> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.417706 mds.mds03 (mds.1
) 14 : cluster [DBG] reconnect by client.37961236 10.141.16.14:0/304725641 after 0.01
-103> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.417815 mds.mds03 (mds.1
) 15 : cluster [DBG] reconnect by client.37960373 10.141.16.18:0/2156995594 after 0.01
-102> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.417935 mds.mds03 (mds.1
) 16 : cluster [DBG] reconnect by client.37960400 10.141.16.16:0/2737464876 after 0.01
-101> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.418627 mds.mds03 (mds.1
) 17 : cluster [DBG] reconnect by client.37865319 10.141.16.21:0/1591580297 after 0.01
-100> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.419295 mds.mds03 (mds.1
) 18 : cluster [DBG] reconnect by client.37865333 10.141.16.20:0/3660809459 after 0.01
-99> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.419728 mds.mds03 (mds.1
) 19 : cluster [DBG] reconnect by client.37865372 10.141.16.19:0/3498217528 after 0.01
-98> 2019-09-19 15:41:50.626 7f036c2f0700 10 log_client logged 2019-09-19 15:41:47.466316 mds.mds03 (mds.1
) 20 : cluster [DBG] reconnect by client.37865356 10.141.16.11:0/1948991158 after 0.0500002
-97> 2019-09-19 15:41:53.336 7f03692ea700 5 mds.beacon.mds03 Sending beacon up:active seq 7
-96> 2019-09-19 15:41:53.336 7f03692ea700 10 monclient: _send_mon_message to mon.mds03 at v2:10.141.16.3:33
00/0
-95> 2019-09-19 15:41:53.346 7f036faf7700 5 mds.beacon.mds03 received beacon reply up:active seq 7 rtt 0.0
1
-94> 2019-09-19 15:41:55.216 7f0369aeb700 2 mds.1.cache Memory usage: total 3260792, rss 2826672, heap 32
3836, baseline 323836, 477 / 888012 inodes have caps, 482 caps, 0.000542785 caps per inode
-93> 2019-09-19 15:41:55.216 7f03652e2700 5 mds.1.log _submit_thread 30520209337247~1811 : EUpdate scatter
_writebehind [metablob 0x10010dcb97f.101*, 2 dirs]
-92> 2019-09-19 15:41:55.216 7f03652e2700 5 mds.1.log _submit_thread 30520209339078~1835 : EUpdate scatter
_writebehind [metablob 0x20005af5c62, 2 dirs]
-91> 2019-09-19 15:41:55.216 7f03652e2700 5 mds.1.log _submit_thread 30520209340933~1531 : EUpdate scatter

```

```

_writebehind [metablob 0x101, 1 dirs]
-90> 2019-09-19 15:41:55.216 7f03652e2700 5 mds.1.log _submit_thread 30520209342484~5656 : EUpdate scatter
_writebehind [metablob 0x20005af5c63, 5 dirs]
-89> 2019-09-19 15:41:55.216 7f03652e2700 5 mds.1.log _submit_thread 30520209348160~2846 : EUpdate scatter
_writebehind [metablob 0x20005af5c6c, 2 dirs]
-88> 2019-09-19 15:41:55.216 7f03652e2700 5 mds.1.log _submit_thread 30520209351026~11258 : EUpdate scatter
r_writebehind [metablob 0x1000bf2626b, 3 dirs]
-87> 2019-09-19 15:41:55.216 7f03652e2700 5 mds.1.log _submit_thread 30520209362304~9900 : EUpdate scatter
_writebehind [metablob 0x1000bf26309, 7 dirs]
-86> 2019-09-19 15:41:55.216 7f03652e2700 5 mds.1.log _submit_thread 30520209372224~8611 : EUpdate scatter
_writebehind [metablob 0x1000bd8ac7b, 2 dirs]
-85> 2019-09-19 15:41:55.216 7f03652e2700 5 mds.1.log _submit_thread 30520209380855~2777 : EUpdate scatter
_writebehind [metablob 0x10010dcb97f.101*, 2 dirs]
-84> 2019-09-19 15:41:55.216 7f03652e2700 5 mds.1.log _submit_thread 30520209383652~2774 : EUpdate scatter
_writebehind [metablob 0x20005af5c62, 2 dirs]
-83> 2019-09-19 15:41:55.216 7f03652e2700 5 mds.1.log _submit_thread 30520209386446~4047 : EUpdate scatter
_writebehind [metablob 0x20005af5c68, 3 dirs]
-82> 2019-09-19 15:41:57.336 7f03692ea700 5 mds.beacon.mds03 Sending beacon up:active seq 8
-81> 2019-09-19 15:41:57.336 7f03692ea700 10 monclient: _send_mon_message to mon.mds03 at v2:10.141.16.3:33
00/0
-80> 2019-09-19 15:41:57.346 7f036faf7700 5 mds.beacon.mds03 received beacon reply up:active seq 8 rtt 0.0
1
-79> 2019-09-19 15:42:00.116 7f0367ae7700 4 mds.1.server handle_client_request client_request(client.37865
319:2375 lookup #0x1001172f39d/koi8-c 2019-09-19 15:41:39.482105 RETRY=1 caller_uid=0, caller_gid=0{0,}) v4
-78> 2019-09-19 15:42:00.116 7f0367ae7700 4 mds.1.server handle_client_request client_request(client.37865
333:8866 lookup #0x100166004d4/Windows-MSVC-C.cmake 2019-09-19 15:41:39.479661 RETRY=1 caller_uid=0, caller_gi
d=0{0,}) v4
-77> 2019-09-19 15:42:00.116 7f0367ae7700 4 mds.1.server handle_client_request client_request(client.37865
372:5807 setattr size=5055 #0x20005a6eb66 2019-09-19 15:41:39.478859 RETRY=1 caller_uid=0, caller_gid=0{0,}) v
4
-76> 2019-09-19 15:42:00.116 7f0367ae7700 4 mds.1.server handle_client_request client_request(client.37865
372:5808 lookup #0x20005a6eb3a/schema.cpython-37.pyc 2019-09-19 15:41:39.589633 RETRY=1 caller_uid=0, caller_g
id=0{0,}) v4
-75> 2019-09-19 15:42:00.116 7f0367ae7700 4 mds.1.server handle_client_request client_request(client.37865
356:3604570 readdir #0x2000090d629 2019-09-19 15:41:39.523863 RETRY=1 caller_uid=0, caller_gid=0{0,}) v4
-74> 2019-09-19 15:42:00.116 7f0367ae7700 5 mds.1.bal mds.1 epoch 9 load mdsload<[pop IRD:[C 3.00e+00] IWR
:[C 7.60e-01] RDR:[C 0.00e+00] FET:[C 0.00e+00] STR:[C 0.00e+00] *LOAD:4.5]
/[pop IRD:[C 3.00e+00] IWR:[C 7.60e-01] RDR:[C 0.00e+00] FET:[C 0.00e+00] STR:[C 0.00e+00] *LOAD:4.5]
, req 0, hr 0, qlen 0, cpu 0>
-73> 2019-09-19 15:42:00.116 7f0367ae7700 5 mds.1.bal import_map from 0 -> 4.52093
-72> 2019-09-19 15:42:00.116 7f0367ae7700 5 mds.1.bal prep_rebalance: cluster loads are
-71> 2019-09-19 15:42:00.116 7f0367ae7700 5 mds.1.bal prep_rebalance: my load 4.52093 target 2.26046
total 4.52093
-70> 2019-09-19 15:42:00.116 7f0367ae7700 5 mds.1.bal mds.0 is underloaded or barely overloaded.
-69> 2019-09-19 15:42:00.116 7f0367ae7700 5 mds.1.bal i am sufficiently overloaded
-68> 2019-09-19 15:42:00.116 7f0367ae7700 5 mds.1.bal - mds.1 exports 2.26046 to mds.0
-67> 2019-09-19 15:42:00.116 7f0367ae7700 5 mds.1.bal rebalance done
-66> 2019-09-19 15:42:00.116 7f03652e2700 5 mds.1.log _submit_thread 30520209390513~25561 : EUpdate setatt
r [metablob 0x1000bd8ac7b, 12 dirs]
-65> 2019-09-19 15:42:00.116 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
372:5809 rmixmap #0x20005a6eb66 security.ima 2019-09-19 15:42:00.124872 caller_uid=0, caller_gid=0{0,}) v4
-64> 2019-09-19 15:42:00.126 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
319:2376 readdir #0x1001172f3d2 2019-09-19 15:42:00.127381 caller_uid=0, caller_gid=0{0,}) v4
-63> 2019-09-19 15:42:00.126 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8867 lookup #0x100166004d4/Windows-MSVC-CXX.cmake 2019-09-19 15:42:00.127537 caller_uid=0, caller_gid=0{0,
}) v4
-62> 2019-09-19 15:42:00.126 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
372:5810 lookup #0x20005af5c74/lxml.cpython-37.pyc 2019-09-19 15:42:00.130435 caller_uid=0, caller_gid=0{0,})
v4
-61> 2019-09-19 15:42:00.126 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8868 lookup #0x100166004d4/Windows-MSVC.cmake 2019-09-19 15:42:00.132050 caller_uid=0, caller_gid=0{0,}) v
4
-60> 2019-09-19 15:42:00.126 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
372:5811 setattr mtime=2019-09-16 12:00:14.047762 atime=2019-09-19 15:42:00.121361 #0x20005a6eb66 2019-09-19 1
5:42:00.135404 caller_uid=0, caller_gid=0{0,}) v4
-59> 2019-09-19 15:42:00.126 7f03652e2700 5 mds.1.log _submit_thread 30520209416094~2363 : EUpdate setattr
[metablob 0x20005af5c6c, 2 dirs]
-58> 2019-09-19 15:42:00.126 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
372:5812 rmixmap #0x20005a6eb66 security.ima 2019-09-19 15:42:00.136273 caller_uid=0, caller_gid=0{0,}) v4
-57> 2019-09-19 15:42:00.126 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8869 lookup #0x100166004d4/Windows-NMcl.cmake 2019-09-19 15:42:00.136439 caller_uid=0, caller_gid=0{0,}) v
4
-56> 2019-09-19 15:42:00.136 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
372:5813 lookup #0x20005af5c74/minidom.cpython-37.pyc 2019-09-19 15:42:00.140199 caller_uid=0, caller_gid=0{0,
}) v4

```

```
-55> 2019-09-19 15:42:00.136 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8870 lookup #0x100166004d4/Windows-NVIDIA-CUDA.cmake 2019-09-19 15:42:00.140572 caller_uid=0, caller_gid=0
{0,}) v4
-54> 2019-09-19 15:42:00.136 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
372:5814 open #0x20005a6eb67 2019-09-19 15:42:00.144319 caller_uid=0, caller_gid=0{0,}) v4
-53> 2019-09-19 15:42:00.136 7f03652e2700 5 mds.1.log _submit_thread 30520209418477~1532 : EOpen [metablob
0x20005af5c74, 1 dirs], 1 open files
-52> 2019-09-19 15:42:00.136 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8871 lookup #0x100166004d4/Windows-OpenWatcom-C.cmake 2019-09-19 15:42:00.144713 caller_uid=0, caller_gid=
0{0,}) v4
-51> 2019-09-19 15:42:00.146 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8872 lookup #0x100166004d4/Windows-OpenWatcom-CXX.cmake 2019-09-19 15:42:00.148593 caller_uid=0, caller_gi
d=0{0,}) v4
-50> 2019-09-19 15:42:00.146 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8873 lookup #0x100166004d4/Windows-OpenWatcom.cmake 2019-09-19 15:42:00.152535 caller_uid=0, caller_gid=0{
0,}) v4
-49> 2019-09-19 15:42:00.156 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8874 lookup #0x100166004d4/Windows-PGI-C.cmake 2019-09-19 15:42:00.156622 caller_uid=0, caller_gid=0{0,})
v4
-48> 2019-09-19 15:42:00.156 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
319:2377 lookup #0x1001172f39d/koi8-r 2019-09-19 15:42:00.158324 caller_uid=0, caller_gid=0{0,}) v4
-47> 2019-09-19 15:42:00.156 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
356:3604571 readdir #0x2000090d62b 2019-09-19 15:42:00.159812 caller_uid=0, caller_gid=0{0,}) v4
-46> 2019-09-19 15:42:00.156 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8875 lookup #0x100166004d4/Windows-PGI-Fortran.cmake 2019-09-19 15:42:00.160603 caller_uid=0, caller_gid=0
{0,}) v4
-45> 2019-09-19 15:42:00.156 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
319:2378 readdir #0x1001172f3d4 2019-09-19 15:42:00.161768 caller_uid=0, caller_gid=0{0,}) v4
-44> 2019-09-19 15:42:00.156 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8876 lookup #0x100166004d4/Windows-PGI.cmake 2019-09-19 15:42:00.164482 caller_uid=0, caller_gid=0{0,}) v4
-43> 2019-09-19 15:42:00.166 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8877 lookup #0x100166004d4/Windows-Watcom-C.cmake 2019-09-19 15:42:00.168639 caller_uid=0, caller_gid=0{0,
}) v4
-42> 2019-09-19 15:42:00.166 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8878 lookup #0x100166004d4/Windows-Watcom-CXX.cmake 2019-09-19 15:42:00.172441 caller_uid=0, caller_gid=0{
0,}) v4
-41> 2019-09-19 15:42:00.166 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8879 lookup #0x100166004d4/Windows-df.cmake 2019-09-19 15:42:00.175790 caller_uid=0, caller_gid=0{0,}) v4
-40> 2019-09-19 15:42:00.176 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8880 lookup #0x100166004d4/Windows-windres.cmake 2019-09-19 15:42:00.179236 caller_uid=0, caller_gid=0{0,}
) v4
-39> 2019-09-19 15:42:00.176 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8881 lookup #0x100166004d4/Windows.cmake 2019-09-19 15:42:00.182733 caller_uid=0, caller_gid=0{0,}) v4
-38> 2019-09-19 15:42:00.176 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8882 lookup #0x100166004d4/WindowsCE-MSVC-C.cmake 2019-09-19 15:42:00.186270 caller_uid=0, caller_gid=0{0,
}) v4
-37> 2019-09-19 15:42:00.186 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
356:3604572 lookup #0x2000090d465/fa 2019-09-19 15:42:00.189287 caller_uid=0, caller_gid=0{0,}) v4
-36> 2019-09-19 15:42:00.186 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
319:2379 lookup #0x1001172f39d/koi8-u 2019-09-19 15:42:00.189393 caller_uid=0, caller_gid=0{0,}) v4
-35> 2019-09-19 15:42:00.186 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8883 lookup #0x100166004d4/WindowsCE-MSVC-CXX.cmake 2019-09-19 15:42:00.189747 caller_uid=0, caller_gid=0{
0,}) v4
-34> 2019-09-19 15:42:00.186 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8884 lookup #0x100166004d4/WindowsCE.cmake 2019-09-19 15:42:00.193093 caller_uid=0, caller_gid=0{0,}) v4
-33> 2019-09-19 15:42:00.186 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
319:2380 readdir #0x1001172f3d5 2019-09-19 15:42:00.193504 caller_uid=0, caller_gid=0{0,}) v4
-32> 2019-09-19 15:42:00.186 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
356:3604573 readdir #0x2000090d49c 2019-09-19 15:42:00.194633 caller_uid=0, caller_gid=0{0,}) v4
-31> 2019-09-19 15:42:00.186 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8885 lookup #0x100166004d4/WindowsPaths.cmake 2019-09-19 15:42:00.196436 caller_uid=0, caller_gid=0{0,}) v
4
-30> 2019-09-19 15:42:00.196 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8886 lookup #0x100166004d4/WindowsPhone-MSVC-C.cmake 2019-09-19 15:42:00.199747 caller_uid=0, caller_gid=0
{0,}) v4
-29> 2019-09-19 15:42:00.196 7f036b2ee700 10 monclient: tick
-28> 2019-09-19 15:42:00.196 7f036b2ee700 10 monclient: _check_auth_rotating have uptodate secrets (they ex
pire after 2019-09-19 15:41:30.200562)
-27> 2019-09-19 15:42:00.196 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8887 lookup #0x100166004d4/WindowsPhone-MSVC-CXX.cmake 2019-09-19 15:42:00.203132 caller_uid=0, caller_gid
=0{0,}) v4
-26> 2019-09-19 15:42:00.196 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
372:5815 lookup #0x20005a6eb3a/selectable.cpython-37.pyc 2019-09-19 15:42:00.204970 caller_uid=0, caller_gid=0
{0,}) v4
-25> 2019-09-19 15:42:00.196 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
```

```

333:8888 lookup #0x100166004d4/WindowsPhone.cmake 2019-09-19 15:42:00.206381 caller_uid=0, caller_gid=0{0,}) v
4
-24> 2019-09-19 15:42:00.206 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8889 lookup #0x100166004d4/WindowsStore-MSVC-C.cmake 2019-09-19 15:42:00.209703 caller_uid=0, caller_gid=0
{0,}) v4
-23> 2019-09-19 15:42:00.206 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8890 lookup #0x100166004d4/WindowsStore-MSVC-CXX.cmake 2019-09-19 15:42:00.213200 caller_uid=0, caller_gid
=0{0,}) v4
-22> 2019-09-19 15:42:00.216 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8891 lookup #0x100166004d4/WindowsStore.cmake 2019-09-19 15:42:00.216577 caller_uid=0, caller_gid=0{0,}) v
4
-21> 2019-09-19 15:42:00.216 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8892 lookup #0x100166004d4/Xenix.cmake 2019-09-19 15:42:00.220230 caller_uid=0, caller_gid=0{0,}) v4
-20> 2019-09-19 15:42:00.216 7f0369aeb700 2 mds.1.cache Memory usage: total 4603496, rss 4167920, heap 32
3836, baseline 323836, 501 / 1162471 inodes have caps, 506 caps, 0.00043528 caps per inode
-19> 2019-09-19 15:42:00.216 7f03652e2700 5 mds.1.log _submit_thread 30520209420029~9062 : EUpdate scatter
_writebehind [metablob 0x1000bd8ac7b, 2 dirs]
-18> 2019-09-19 15:42:00.216 7f03652e2700 5 mds.1.log _submit_thread 30520209429111~10579 : EUpdate scatte
r_writebehind [metablob 0x1000bf26309, 9 dirs]
-17> 2019-09-19 15:42:00.216 7f03652e2700 5 mds.1.log _submit_thread 30520209439710~2305 : EUpdate scatter
_writebehind [metablob 0x1000bf2745b.001*, 2 dirs]
-16> 2019-09-19 15:42:00.216 7f03652e2700 5 mds.1.log _submit_thread 30520209442035~1845 : EUpdate scatter
_writebehind [metablob 0x1000c233753, 2 dirs]
-15> 2019-09-19 15:42:00.216 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8893 lookup #0x100166004d4/eCos.cmake 2019-09-19 15:42:00.223360 caller_uid=0, caller_gid=0{0,}) v4
-14> 2019-09-19 15:42:00.216 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
319:2381 lookup #0x1001172f39d/microsoft-cp1251 2019-09-19 15:42:00.224940 caller_uid=0, caller_gid=0{0,}) v4
-13> 2019-09-19 15:42:00.226 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8894 lookup #0x100166004d4/gas.cmake 2019-09-19 15:42:00.226624 caller_uid=0, caller_gid=0{0,}) v4
-12> 2019-09-19 15:42:00.226 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
319:2382 readdir #0x1001172f3a7 2019-09-19 15:42:00.228673 caller_uid=0, caller_gid=0{0,}) v4
-11> 2019-09-19 15:42:00.226 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8895 lookup #0x100166004d4/kFreeBSD.cmake 2019-09-19 15:42:00.229668 caller_uid=0, caller_gid=0{0,}) v4
-10> 2019-09-19 15:42:00.226 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8896 lookup #0x100166004d4/syllable.cmake 2019-09-19 15:42:00.232746 caller_uid=0, caller_gid=0{0,}) v4
-9> 2019-09-19 15:42:00.236 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
333:8897 readdir #0x10016601379 2019-09-19 15:42:00.240672 caller_uid=0, caller_gid=0{0,}) v4
-8> 2019-09-19 15:42:00.236 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
356:3604574 readdir #0x2000090d630 2019-09-19 15:42:00.241832 caller_uid=0, caller_gid=0{0,}) v4
-7> 2019-09-19 15:42:00.266 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
356:3604575 readdir #0x2000090d631 2019-09-19 15:42:00.272158 caller_uid=0, caller_gid=0{0,}) v4
-6> 2019-09-19 15:42:00.326 7f03652e2700 5 mds.1.log _submit_thread 30520209443900~3089 : EUpdate scatter
_writebehind [metablob 0x20005af5c63, 3 dirs]
-5> 2019-09-19 15:42:00.326 7f03652e2700 5 mds.1.log _submit_thread 30520209447009~10579 : EUpdate scatte
r_writebehind [metablob 0x1000bf26309, 9 dirs]
-4> 2019-09-19 15:42:00.326 7f03652e2700 5 mds.1.log _submit_thread 30520209457608~2305 : EUpdate scatter
_writebehind [metablob 0x1000bf2745b.001*, 2 dirs]
-3> 2019-09-19 15:42:00.326 7f03652e2700 5 mds.1.log _submit_thread 30520209459933~1030 : EUpdate check_i
node_max_size [metablob 0x20005af5c74, 1 dirs]
-2> 2019-09-19 15:42:00.326 7f036c2f0700 4 mds.1.server handle_client_request client_request(client.37865
372:5816 setattr size=1138 #0x20005a6eb67 2019-09-19 15:42:00.333015 caller_uid=0, caller_gid=0{0,}) v4
-1> 2019-09-19 15:42:00.336 7f036c2f0700 -1 /home/jenkins-build/build/workspace/ceph-build/ARCH/x86_64/AVA
ILABLE_ARCH/x86_64/AVAILABLE_DIST/centos7/DIST/centos7/MACHINE_SIZE/huge/release/14.2.3/rpm/el7/BUILD/ceph-14.
2.3/src/mds/MDCache.cc: In function 'CInode* MDCache::cow_inode(CInode*, snapid_t)' thread 7f036c2f0700 time 2
019-09-19 15:42:00.333567
/home/jenkins-build/build/workspace/ceph-build/ARCH/x86_64/AVAILABLE_ARCH/x86_64/AVAILABLE_DIST/centos7/DIST/c
entos7/MACHINE_SIZE/huge/release/14.2.3/rpm/el7/BUILD/ceph-14.2.3/src/mds/MDCache.cc: 1498: FAILED ceph_assert
(!lock->get_num_wrllocks())

```

```

ceph version 14.2.3 (0f776cf838alae3130b2b73dc26be9c95c6ccc39) nautilus (stable)
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x14a) [0x7f0375773ac2]
2: (ceph::__ceph_assertf_fail(char const*, char const*, int, char const*, char const*, ...) +0) [0x7f0375773c9
0]
3: ((+0x1ee48d) [0x55a7e4ccf48d]
4: (MDCache::journal_cow_dentry(MutationImpl*, EMetaBlob*, CDentry*, snapid_t, CInode**, CDentry::linkage_t*)
+0x823) [0x55a7e4ccfb3]
5: (MDCache::journal_dirty_inode(MutationImpl*, EMetaBlob*, CInode*, snapid_t)+0xbc) [0x55a7e4cd042c]
6: (Locker::_do_cap_update(CInode*, Capability*, int, snapid_t, boost::intrusive_ptr<MClientCaps const> const
&, boost::intrusive_ptr<MClientCaps> const&, bool*)+0xfb6) [0x55a7e4d957e6]
7: (Locker::handle_client_caps(boost::intrusive_ptr<MClientCaps const> const&)+0x2059) [0x55a7e4d9c8e9]
8: (Locker::dispatch(boost::intrusive_ptr<Message const> const&)+0xe7) [0x55a7e4daaf97]
9: (MDSRank::handle_deferrable_message(boost::intrusive_ptr<Message const> const&)+0x304) [0x55a7e4c089e4]
10: (MDSRank::_dispatch(boost::intrusive_ptr<Message const> const&, bool)+0x6eb) [0x55a7e4c0b1bb]
11: (MDSRankDispatcher::ms_dispatch(boost::intrusive_ptr<Message const> const&)+0x40) [0x55a7e4c0b8d0]
12: (MDSDaemon::ms_dispatch2(boost::intrusive_ptr<Message const> const&)+0xfc) [0x55a7e4bf82ac]

```

```
13: (DispatchQueue::entry()+0x12a9) [0x7f0375968dd9]
14: (DispatchQueue::DispatchThread::entry()+0xd) [0x7f0375a183dd]
15: (()+0x7dd5) [0x7f0373642dd5]
16: (clone()+0x6d) [0x7f03722f302d]
```

```
0> 2019-09-19 15:42:00.336 7f036c2f0700 -1 *** Caught signal (Aborted) **
in thread 7f036c2f0700 thread_name:ms_dispatch
```

```
ceph version 14.2.3 (0f776cf838a1ae3130b2b73dc26be9c95c6ccc39) nautilus (stable)
```

```
1: (()+0xf5d0) [0x7f037364a5d0]
2: (gsignal()+0x37) [0x7f037222b2c7]
3: (abort()+0x148) [0x7f037222c9b8]
4: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x199) [0x7f0375773b11]
5: (ceph::__ceph_assertf_fail(char const*, char const*, int, char const*, char const*, ...) +0) [0x7f0375773c90]
6: (()+0x1ee48d) [0x55a7e4ccf48d]
7: (MDCache::journal_cow_dentry(MutationImpl*, EMetaBlob*, CDentry*, snapid_t, CInode**, CDentry::linkage_t*)+0x823) [0x55a7e4ccfcb3]
8: (MDCache::journal_dirty_inode(MutationImpl*, EMetaBlob*, CInode*, snapid_t)+0xbc) [0x55a7e4cd042c]
9: (Locker::_do_cap_update(CInode*, Capability*, int, snapid_t, boost::intrusive_ptr<MClientCaps const> const&, boost::intrusive_ptr<MClientCaps> const&, bool*)+0xf6) [0x55a7e4d957e6]
10: (Locker::handle_client_caps(boost::intrusive_ptr<MClientCaps const> const&)+0x2059) [0x55a7e4d9c8e9]
11: (Locker::dispatch(boost::intrusive_ptr<Message const> const&)+0xe7) [0x55a7e4daaf97]
12: (MDSRank::handle_deferrable_message(boost::intrusive_ptr<Message const> const&)+0x304) [0x55a7e4c089e4]
13: (MDSRank::_dispatch(boost::intrusive_ptr<Message const> const&, bool)+0x6eb) [0x55a7e4c0b1bb]
14: (MDSRankDispatcher::ms_dispatch(boost::intrusive_ptr<Message const> const&)+0x40) [0x55a7e4c0b8d0]
15: (MDSDaemon::ms_dispatch2(boost::intrusive_ptr<Message> const&)+0xfc) [0x55a7e4bf82ac]
16: (DispatchQueue::entry()+0x12a9) [0x7f0375968dd9]
17: (DispatchQueue::DispatchThread::entry()+0xd) [0x7f0375a183dd]
18: (()+0x7dd5) [0x7f0373642dd5]
19: (clone()+0x6d) [0x7f03722f302d]
NOTE: a copy of the executable, or `objdump -rds <executable>` is needed to interpret this.
```

```
--- logging levels ---
```

```
0/ 5 none
0/ 1 lockdep
0/ 1 context
1/ 1 crush
1/ 5 mds
1/ 5 mds_balancer
1/ 5 mds_locker
1/ 5 mds_log
1/ 5 mds_log_expire
1/ 5 mds_migrator
0/ 1 buffer
0/ 1 timer
0/ 1 filer
0/ 1 striper
0/ 1 objecter
0/ 5 rados
0/ 5 rbd
0/ 5 rbd_mirror
0/ 5 rbd_replay
0/ 5 journaler
0/ 5 objectcacher
0/ 5 client
1/ 5 osd
0/ 5 optracker
0/ 5 objclass
1/ 3 filestore
1/ 3 journal
0/ 0 ms
1/ 5 mon
0/10 monc
1/ 5 paxos
0/ 5 tp
1/ 5 auth
1/ 5 crypto
1/ 1 finisher
1/ 1 reserver
1/ 5 heartbeatmap
1/ 5 perfcounter
1/ 5 rgw
1/ 5 rgw_sync
1/10 civetweb
```

```
1/ 5 javaclient
1/ 5 asok
1/ 1 throttle
0/ 0 refs
1/ 5 xio
1/ 5 compressor
1/ 5 bluestore
1/ 5 bluefs
1/ 3 bdev
1/ 5 kstore
4/ 5 rocksdb
4/ 5 leveldb
4/ 5 memdb
1/ 5 kinetic
1/ 5 fuse
1/ 5 mgr
1/ 5 mgrc
1/ 5 dpdk
1/ 5 eventtrace
-2/-2 (syslog threshold)
-1/-1 (stderr threshold)
max_recent      10000
max_new         1000
log_file /var/log/ceph/ceph-mds.mds03.log
--- end dump of recent events ---
[root@mds03 ~]#
```

#2 - 09/19/2019 03:42 PM - Dan van der Ster

Looks like the backport of <https://tracker.ceph.com/issues/39987> to nautilus was incomplete, it's missing <https://github.com/ceph/ceph/pull/28459>

#3 - 09/22/2019 07:34 AM - Zheng Yan

<https://tracker.ceph.com/issues/41948>

#4 - 09/23/2019 02:11 AM - Patrick Donnelly

- Status changed from New to Duplicate

#5 - 09/23/2019 02:12 AM - Patrick Donnelly

- Duplicates Bug #41948: nautilus: mds: incomplete backport of #40444 (MDCache::cow_inode does not cleanup unneeded client_snap_caps) added