

## fs - Bug #42467

### mds: daemon crashes while updating blacklist

10/24/2019 12:31 PM - Ondřej Kopka

<b>Status:</b>	Duplicate	<b>% Done:</b>	0%
<b>Priority:</b>	High		
<b>Assignee:</b>	Zheng Yan		
<b>Category:</b>			
<b>Target version:</b>			
<b>Source:</b>	Community (user)	<b>Affected Versions:</b>	v13.2.8, v14.2.5
<b>Tags:</b>		<b>ceph-qa-suite:</b>	
<b>Backport:</b>		<b>Component(FS):</b>	MDS
<b>Regression:</b>	No	<b>Labels (FS):</b>	crash
<b>Severity:</b>	3 - minor	<b>Pull request ID:</b>	33570
<b>Reviewed:</b>		<b>Crash signature:</b>	

#### Description

Ubuntu 18.04.3 LTS

ceph version 14.2.4 (75f4de193b3ea58512f204623e6c5a16e6c1e1ba) nautilus (stable)

We have setup with 2 MDS in standby replay mode.

Can you please look into this issue? Let me know if you will want me to provide more information.

Our MDS daemon sometimes crashes with following error:

```
-31> 2019-10-23 06:33:12.831 7fbd4af2e700 10 monclient: tick
-30> 2019-10-23 06:33:12.831 7fbd4af2e700 10 monclient: _check_auth_rotating have uptodate secrets (they expire after 2019-10-23 06:32:42.838088)
-29> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client log_queue is 14 last_log 1052 sent 1038 num 14 unsent 14 sending 14
-28> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client will send 2019-10-23 06:33:07.230391 mds.ceph-mds-2.domain.com (mds.0) 1039 : cluster [WRN] 67 slow requests, 3 included below; oldest blocked for > 1605.678039 secs
-27> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client will send 2019-10-23 06:33:07.230402 mds.ceph-mds-2.domain.com (mds.0) 1040 : cluster [WRN] slow request 483.710501 seconds old, received at 2019-10-23 06:25:03.519777: client_request(client.1614065:109 getattr AsLsXsFs #0x100094b06b5 2019-10-23 06:25:02.988000 caller_uid=80109, caller_gid=90404{}) currently cleaned up request
-26> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client will send 2019-10-23 06:33:07.230410 mds.ceph-mds-2.domain.com (mds.0) 1041 : cluster [WRN] slow request 481.450746 seconds old, received at 2019-10-23 06:25:05.779532: client_request(client.1614065:1747 create #0x100094b06b5/.pureftpd-upload.5daff23f.15.829.b9ff6bce 2019-10-23 06:25:05.248000 caller_uid=80109, caller_gid=90404{}) currently cleaned up request
-25> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client will send 2019-10-23 06:33:07.230416 mds.ceph-mds-2.domain.com (mds.0) 1042 : cluster [WRN] slow request 484.701491 seconds old, received at 2019-10-23 06:25:02.528787: client_request(client.1614065:55 mkdir #0x10005768670/443552 2019-10-23 06:25:01.996000 caller_uid=80213, caller_gid=91584{}) currently cleaned up request
-24> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client will send 2019-10-23 06:33:12.230570 mds.ceph-mds-2.domain.com (mds.0) 1043 : cluster [WRN] 68 slow requests, 5 included below; oldest blocked for > 1610.678190 secs
-23> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client will send 2019-10-23 06:33:12.230580 mds.ceph-mds-2.domain.com (mds.0) 1044 : cluster [WRN] slow request 124.476924 seconds old, received at 2019-10-23 06:31:07.753505: client_request(client.1762020:72 getattr AsLsXsFs #0x100094b06b5 2019-10-23 06:31:06.964000 caller_uid=80109, caller_gid=90404{}) currently cleaned up request
-22> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client will send 2019-10-23 06:33:12.230588 mds.ceph-mds-2.domain.com (mds.0) 1045 : cluster [WRN] slow request 64.468477 seconds old, received at 2019-10-23 06:32:07.761952: client_request(client.1647734:93 lookup #0x100096dcc7e/.pureftpd-upload.5daff2dd.15.6f5.5f36917c 2019-10-23 06:32:06.896000 caller_uid=82183, caller_gid=96143{}) currently failed to rdlock, waiting
```

```

-21> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client will send 2019-10-23 06:33:12.230594 m
ds.ceph-mds-2.domain.com (mds.0) 1046 : cluster [WRN] slow request 34.478042 seconds old, received
at 2019-10-23 06:32:37.752387: client_request(client.1669957:490 create #0x1/watchdog-rnd-ceph-ht
tp-1-89-187-188-153 2019-10-23 06:32:37.744276 caller_uid=0, caller_gid=0{ }) currently failed to x
lock, waiting
-20> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client will send 2019-10-23 06:33:12.230603 m
ds.ceph-mds-2.domain.com (mds.0) 1047 : cluster [WRN] slow request 244.119127 seconds old, receive
d at 2019-10-23 06:29:08.111302: client_request(client.1669922:50 getattr AsLsXsFs #0x100094b06b5
2019-10-23 06:29:07.804000 caller_uid=80109, caller_gid=90404{ }) currently cleaned up request
-19> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client will send 2019-10-23 06:33:12.230609 m
ds.ceph-mds-2.domain.com (mds.0) 1048 : cluster [WRN] slow request 244.554605 seconds old, receive
d at 2019-10-23 06:29:07.675824: client_request(client.1669922:49 create #0x100094b06b5/.pureftpd-
upload.5daff331.15.854.3e0f0ef8 2019-10-23 06:29:07.368000 caller_uid=80109, caller_gid=90404{ }) c
urrently cleaned up request
-18> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client will send 2019-10-23 06:33:12.285084 m
ds.ceph-mds-2.domain.com (mds.0) 1049 : cluster [WRN] evicting unresponsive client ceph-http-1-89-
187-188-153.domain.com:http-1 (1614114), after 68.7832 seconds
-17> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client will send 2019-10-23 06:33:12.285097 m
ds.ceph-mds-2.domain.com (mds.0) 1050 : cluster [INF] Evicting (and blacklisting) client session 1
614114 (v1:X.X.X.X:0/1386239241)
-16> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client will send 2019-10-23 06:33:12.285130 m
ds.ceph-mds-2.domain.com (mds.0) 1051 : cluster [WRN] evicting unresponsive client push-zone-push-
33-b-ceph-89-187-188-155.domain.com:push-33b (1647734), after 63.0713 seconds
-15> 2019-10-23 06:33:12.831 7fbd4af2e700 10 log_client will send 2019-10-23 06:33:12.285136 m
ds.ceph-mds-2.domain.com (mds.0) 1052 : cluster [INF] Evicting (and blacklisting) client session 1
647734 (v1:X.X.X.X:0/2412026873)
-14> 2019-10-23 06:33:12.831 7fbd4af2e700 10 monclient: _send_mon_message to mon.ceph-osd-3 at
v2:X.X.X.X:3300/0
-13> 2019-10-23 06:33:12.839 7fbd4bf30700 4 mds.0.server handle_client_request client_request(
client.1651513:1510574 open #0x10002e2cf9d/How.to.Get.Away.with.Murder.S06E03.mp4 2019-10-23 06:33
:12.839252 caller_uid=0, caller_gid=0{ }) v2
-12> 2019-10-23 06:33:12.839 7fbd44f22700 5 mds.0.log _submit_thread 3890730238094~1011 : EOpe
n [metablob 0x10002e2cf9d.101*, 1 dirs], 1 open files
-11> 2019-10-23 06:33:12.859 7fbd44f22700 5 mds.0.log _submit_thread 3890730239125~988 : EOpen
[metablob 0x10009b2c387, 1 dirs], 1 open files
-10> 2019-10-23 06:33:12.863 7fbd44f22700 5 mds.0.log _submit_thread 3890730240133~987 : EOpen
[metablob 0x100055504d2, 1 dirs], 1 open files
-9> 2019-10-23 06:33:12.879 7fbd4bf30700 4 mds.0.server handle_client_request client_request(
client.1651513:1510575 lookup #0x100065f2d53/b778efb6098670.ts 2019-10-23 06:33:12.879252 caller_u
id=0, caller_gid=0{ }) v2
-8> 2019-10-23 06:33:12.887 7fbd4bf30700 4 mds.0.server handle_client_request client_request(
client.1669978:331 lookup #0x10004799af8/artifacts.jar 2019-10-23 06:33:12.889642 caller_uid=0, ca
lller_gid=0{ }) v2
-7> 2019-10-23 06:33:12.903 7fbd44f22700 5 mds.0.log _submit_thread 3890730241140~990 : EOpen
[metablob 0x100065f2d53, 1 dirs], 1 open files
-6> 2019-10-23 06:33:12.919 7fbd4bf30700 4 mds.0.server handle_client_request client_request(
client.1651513:1510576 lookup #0x100065f1608/Y70PD6U6817078.ts 2019-10-23 06:33:12.919252 caller_u
id=0, caller_gid=0{ }) v2
-5> 2019-10-23 06:33:12.919 7fbd4bf30700 4 mds.0.server handle_client_request client_request(
client.1651513:1510577 lookup #0x100065f2b6c/5n3Xr08k34552815.ts 2019-10-23 06:33:12.919252 caller
_uid=0, caller_gid=0{ }) v2
-4> 2019-10-23 06:33:12.923 7fbd4bf30700 4 mds.0.161128 handle_osd_map epoch 24683, 2 new bla
cklist entries
-3> 2019-10-23 06:33:12.923 7fbd44f22700 5 mds.0.log _submit_thread 3890730242150~107 : ESess
ion client.1614114 v1:X.X.X.X:0/1386239241 close cmapv 163844044
-2> 2019-10-23 06:33:12.923 7fbd44f22700 5 mds.0.log _submit_thread 3890730242277~171 : ESess
ion client.1647734 v1:X.X.X.X:0/2412026873 close cmapv 163844045 (1000 inos, v350819)
-1> 2019-10-23 06:33:12.927 7fbd4bf30700 -1 /build/ceph-14.2.4/src/mds/MDCache.cc: In function
'MDRequestRef MDCache::request_get(metareqid_t)' thread 7fbd4bf30700 time 2019-10-23 06:33:12.931
234
/build/ceph-14.2.4/src/mds/MDCache.cc: 9306: FAILED ceph_assert(p != active_requests.end())

ceph version 14.2.4 (75f4de193b3ea58512f204623e6c5a16e6c1e1ba) nautilus (stable)
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x152) [0x7fbd54fbeb66]
2: (()+0x27ad3e) [0x7fbd54fbed3e]
3: (MDCache::request_get(metareqid_t)+0xcc) [0x55987812f5bc]
4: (Server::journal_close_session(Session*, int, Context*)+0xblf) [0x5598780826ef]

```

```
5: (Server::kill_session(Session*, Context*)+0x234) [0x559878085ad4]
6: (Server::apply_blacklist(std::set<entity_addr_t, std::less<entity_addr_t>, std::allocator<entity_addr_t> > const&)+0x14d) [0x55987808881d]
7: (MDSRankDispatcher::handle_osd_map()+0x230) [0x559878055e60]
8: (MDSDaemon::handle_core_message(boost::intrusive_ptr<Message const> const&)+0x39e) [0x559878039f0e]
9: (MDSDaemon::ms_dispatch2(boost::intrusive_ptr<Message> const&)+0xc3) [0x55987803a693]
10: (DispatchQueue::entry()+0x1a49) [0x7fbd551f53b9]
11: (DispatchQueue::DispatchThread::entry()+0xd) [0x7fbd552a028d]
12: (()+0x76db) [0x7fbd548bc6db]
13: (clone()+0x3f) [0x7fbd53aa288f]
```

```
0> 2019-10-23 06:33:12.931 7fbd4bf30700 -1 *** Caught signal (Aborted) **
in thread 7fbd4bf30700 thread_name:ms_dispatch
```

```
ceph version 14.2.4 (75f4de193b3ea58512f204623e6c5a16e6c1e1ba) nautilus (stable)
```

```
1: (()+0x12890) [0x7fbd548c7890]
2: (gsignal()+0xc7) [0x7fbd539bfe97]
3: (abort()+0x141) [0x7fbd539c1801]
4: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x1a3) [0x7fbd54fbbeb7]
5: (()+0x27ad3e) [0x7fbd54fbed3e]
6: (MDCache::request_get(metareqid_t)+0xcc) [0x55987812f5bc]
7: (Server::journal_close_session(Session*, int, Context*)+0xb1f) [0x5598780826ef]
8: (Server::kill_session(Session*, Context*)+0x234) [0x559878085ad4]
9: (Server::apply_blacklist(std::set<entity_addr_t, std::less<entity_addr_t>, std::allocator<entity_addr_t> > const&)+0x14d) [0x55987808881d]
10: (MDSRankDispatcher::handle_osd_map()+0x230) [0x559878055e60]
11: (MDSDaemon::handle_core_message(boost::intrusive_ptr<Message const> const&)+0x39e) [0x559878039f0e]
12: (MDSDaemon::ms_dispatch2(boost::intrusive_ptr<Message> const&)+0xc3) [0x55987803a693]
13: (DispatchQueue::entry()+0x1a49) [0x7fbd551f53b9]
14: (DispatchQueue::DispatchThread::entry()+0xd) [0x7fbd552a028d]
15: (()+0x76db) [0x7fbd548bc6db]
16: (clone()+0x3f) [0x7fbd53aa288f]
```

```
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 perfcouter
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.ceph-mds-2.domain.com.log
--- end dump of recent events ---
```

## History

---

### #1 - 10/28/2019 09:06 PM - Greg Farnum

- Project changed from Ceph to fs

### #2 - 10/28/2019 10:00 PM - Patrick Donnelly

- Description updated  
- Target version deleted (v14.2.5)  
- Start date deleted (10/24/2019)  
- Source set to Community (user)  
- Affected Versions v14.2.5 added  
- Component(FS) MDS added  
- Labels (FS) crash added

### #3 - 10/29/2019 01:52 PM - Patrick Donnelly

- Subject changed from MDS daemon crashes to mds: daemon crashes while updating blacklist  
- Assignee set to Patrick Donnelly

### #4 - 01/15/2020 07:39 PM - Patrick Donnelly

- Assignee changed from Patrick Donnelly to Zheng Yan  
- Target version set to v15.0.0  
- Backport set to nautilus,mimic

Zheng, I think you may have inadvertently fixed this in

```
commit a10f051af25640def2ca6132de42693975de5ba1 HEAD
Author: Yan, Zheng <zyan@redhat.com>
Date: Mon Sep 30 14:20:17 2019 +0800
```

```
    mds: cleanup code that kills session requests

    Signed-off-by: "Yan, Zheng" <zyan@redhat.com>
---
src/mds/Server.cc | 7 ++-----
1 file changed, 2 insertions(+), 5 deletions(-)

diff --git a/src/mds/Server.cc b/src/mds/Server.cc
index af48d93a3ec..4012a13b2bd 100644
--- a/src/mds/Server.cc
+++ b/src/mds/Server.cc
@@ -1234,11 +1234,8 @@ void Server::journal_close_session(Session *session, int state, Context *on_safe
     mdlog->flush();

    // clean up requests, too
-   elist<MDRequestImpl*>::iterator p =
-     session->requests.begin(member_offset(MDRequestImpl,
-                                             item_session_request));
-   while (!p.end()) {
-     MDRequestRef mdr = mdcache->request_get((*p)->reqid);
+   for (auto p = session->requests.begin(); !p.end(); ) {
+     MDRequestRef mdr(*p);
+     ++p;
+     mdcache->request_kill(mdr);
  }
}
```

But I'm not sure if it's normal for a request to not be in the `MDCache::active\_requests`. Does it make sense?

## #5 - 01/21/2020 03:04 PM - Zheng Yan

No idea how this happen. when a request is removed from active\_requests, request should also be removed from session->requests

## #6 - 01/21/2020 05:06 PM - Patrick Donnelly

- Status changed from New to Can't reproduce
- Target version deleted (v15.0.0)
- Backport deleted (nautilus,mimic)

## #7 - 02/13/2020 05:30 PM - Patrick Donnelly

- Status changed from Can't reproduce to New
- Priority changed from Normal to High
- Affected Versions v13.2.8 added

A new report of this coming from ceph-users:

"[ceph-users] Ceph MDS ASSERT In function 'MDRequestRef'"

Different trace:

```
-10001> 2020-02-13 17:42:35.543 7f11b5669700 -1 /build/ceph-13.2.8/src/mds/MDCache.cc: In function 'MDRequestRef MDCache::request_get(metareqid_t)' thread 7f11b5669700 time 2020-02-13 17:42:35.545815
/build/ceph-13.2.8/src/mds/MDCache.cc: 9523: FAILED assert(p != active_requests.end())

ceph version 13.2.8 (5579a94fafbc1f9cc913a0f5d362953a5d9c3ae0) mimic (stable)
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x14e) [0x7f11bd8e69de]
2: (()+0x287b67) [0x7f11bd8e6b67]
3: (MDCache::request_get(metareqid_t)+0x94) [0x560cde8bb214]
4: (Server::journal_close_session(Session*, int, Context*)+0x9dd) [0x560cde829d1d]
5: (Server::handle_client_session(MClientSession*)+0x1071) [0x560cde82b0f1]
6: (Server::dispatch(Message*)+0x30b) [0x560cde86f87b]
7: (MDSRank::handle_deferrable_message(Message*)+0x434) [0x560cde7e1664]
8: (MDSRank::_dispatch(Message*, bool)+0x89b) [0x560cde7f8c7b]
9: (MDSRankDispatcher::ms_dispatch(Message*)+0xa3) [0x560cde7f92e3]
10: (MDSDaemon::ms_dispatch(Message*)+0xd3) [0x560cde7d92b3]
11: (DispatchQueue::entry()+0xb92) [0x7f11bd9a9e52]
12: (DispatchQueue::DispatchThread::entry()+0xd) [0x7f11bda46e2d]
13: (()+0x76db) [0x7f11bd1d76db]
14: (clone()+0x3f) [0x7f11bc3bd88f]
```

**#8 - 02/13/2020 10:35 PM - Stefan Kooman**

ceph-post-file: 44655e58-39e0-4fff-a2fc-2645b131c594 for the crash listed above (13.2.8)

**#9 - 03/23/2020 12:16 PM - Zheng Yan**

- Status changed from New to Fix Under Review

- Pull request ID set to 33570

**#10 - 03/30/2020 09:09 AM - Zheng Yan**

- Status changed from Fix Under Review to Duplicate

dup of [#44316](#)