

CephFS - Bug #23812

mds: may send LOCK_SYNC_MIX message to starting MDS

04/21/2018 12:28 AM - Patrick Donnelly

Status:	Resolved	% Done:	0%
Priority:	Urgent		
Assignee:	Zheng Yan		
Category:	Correctness/Safety		
Target version:	v13.0.0		
Source:	Development	Affected Versions:	
Tags:		ceph-qa-suite:	
Backport:	luminous	Component(FS):	MDS
Regression:	No	Labels (FS):	multimds
Severity:	2 - major	Pull request ID:	
Reviewed:		Crash signature:	

Description

From mds.0:

```
2018-04-20 20:01:26.892 7ff249f42700 1 -- 127.0.0.1:6829/4093013988 _send_message--> mds.1 127.0.0.1:6827/3615748515 -- mdsmap(e 32) v1 -- ?+0 0x2e85f98080
2018-04-20 20:01:26.892 7ff249f42700 1 -- 127.0.0.1:6829/4093013988 --> 127.0.0.1:6827/3615748515 -- mdsmap(e 32) v1 -- 0x2e85f98080 con 0
2018-04-20 20:01:26.892 7ff249f42700 1 -- 127.0.0.1:6829/4093013988 _send_message--> mds.1 127.0.0.1:6827/3615748515 -- lock(a=mix inest 0x1.head) v1 -- ?+0 0x2e7d789440
2018-04-20 20:01:26.892 7ff249f42700 1 -- 127.0.0.1:6829/4093013988 --> 127.0.0.1:6827/3615748515 -- lock(a=mix inest 0x1.head) v1 -- 0x2e7d789440 con 0
```

mds.1:

```
2018-04-20 20:01:26.896 7f018cd36700 1 -- 127.0.0.1:6827/3615748515 <== mds.0 127.0.0.1:6829/4093013988 2 ==== mdsmap(e 32) v1 ==== 780+0+0 (4159823880 0 0) 0x3209318a80 con 0x3209430e00
2018-04-20 20:01:26.896 7f018cd36700 5 mds.a handle_mds_map epoch 32 from mds.0
2018-04-20 20:01:26.896 7f018cd36700 5 mds.a old map epoch 32 <= 32, discarding
2018-04-20 20:01:26.896 7f018cd36700 1 -- 127.0.0.1:6827/3615748515 <== mds.0 127.0.0.1:6829/4093013988 3 ==== lock(a=mix inest 0x1.head) v1 ==== 291+0+0 (4212595484 0 0) 0x32091a5e40 con 0x3209430e00
2018-04-20 20:01:26.896 7f018cd36700 -1 /home/pdonnell/ceph/src/mds/Locker.cc: In function 'void Locker::handle_lock(MLock*)' thread 7f018cd36700 time 2018-04-20 20:01:26.898953 /home/pdonnell/ceph/src/mds/Locker.cc: 3870: FAILED assert(mds->is_rejoin() || mds->is_clientreply() || mds->is_active() || mds->is_stopping())
```

```
ceph version 13.0.2-1597-g94271de (94271de7ff6ed4f05c9415cf81e493677adble6d) mimic (dev)
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x102) [0x7f0194551e92]
2: (()+0x298067) [0x7f0194552067]
3: (Locker::handle_lock(MLock*)+0x1c0) [0x32078255b0]
4: (MDSRank::handle_deferrable_message(Message*)+0x545) [0x32076bb5b5]
5: (MDSRank::_dispatch(Message*, bool)+0x62b) [0x32076c71fb]
6: (MDSRankDispatcher::ms_dispatch(Message*)+0x15) [0x32076c78f5]
7: (MDSDaemon::ms_dispatch(Message*)+0xd3) [0x32076b36b3]
8: (DispatchQueue::entry()+0xb5a) [0x7f01945cbaba]
9: (DispatchQueue::DispatchThread::entry()+0xd) [0x7f019466b5cd]
10: (()+0x76ba) [0x7f0193e336ba]
11: (clone()+0x6d) [0x7f01930bd41d]
NOTE: a copy of the executable, or `objdump -rds <executable>` is needed to interpret this.
```

This condition looks wrong:

<https://github.com/ceph/ceph/blob/6f60a995de5645667f2c330d03459a7a9ca469f9/src/mds/Locker.cc#L893-L894>

Related issues:

Related to CephFS - Bug #23814: mds: newly active mds aborts may abort in han...	Rejected	04/21/2018
Copied to CephFS - Backport #23935: luminous: mds: may send LOCK_SYNC_MIX mes...	Resolved	

History

#1 - 04/21/2018 12:30 AM - Patrick Donnelly

- Status changed from New to Fix Under Review

<https://github.com/ceph/ceph/pull/21577>

#2 - 04/23/2018 08:37 AM - Zheng Yan

- Related to Bug #23814: mds: newly active mds aborts may abort in handle_file_lock added

#3 - 04/23/2018 05:50 PM - Patrick Donnelly

- Assignee changed from Patrick Donnelly to Zheng Yan

<https://github.com/ceph/ceph/pull/21601>

#4 - 04/29/2018 08:06 PM - Patrick Donnelly

- Status changed from Fix Under Review to Pending Backport

#5 - 04/29/2018 08:30 PM - Nathan Cutler

- Copied to Backport #23935: luminous: mds: may send LOCK_SYNC_MIX message to starting MDS added

#6 - 06/09/2018 11:15 AM - Nathan Cutler

- Status changed from Pending Backport to Resolved