

fs - Bug #26869

mds: becomes stuck in up:replay during thrashing (not multimds)

08/06/2018 06:35 PM - Patrick Donnelly

Status:	New	Start date:	08/06/2018
Priority:	High	Due date:	
Assignee:	Patrick Donnelly	% Done:	0%
Category:	Correctness/Safety	Estimated time:	0.00 hour
Target version:	v15.0.0	Affected Versions:	
Source:	Q/A	ceph-qa-suite:	
Tags:		Component(FS):	MDS
Backport:	mimic,luminous	Labels (FS):	task(medium)
Regression:	No	Pull request ID:	
Severity:	3 - minor		
Reviewed:			

Description

```
2018-08-05 16:23:02.504 7fd0fcd23700 10 mds.a map says I am 172.21.15.26:6817/673291871 mds.0.14 s
tate up:replay
2018-08-05 16:23:02.504 7fd0fcd23700 4 mds.0.purge_queue operator(): data pool 3 not found in OS
DMap
2018-08-05 16:23:02.505 7fd0fcd23700 10 mds.a handle_mds_map: initializing MDS rank 0
2018-08-05 16:23:02.505 7fd0fcd23700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.26:6789/0 --
mon_subscribe({osdmap=0}) v3 -- 0x56359304f400 con 0
2018-08-05 16:23:02.505 7fd0fcd23700 10 mds.0.0 update_log_config log_to_monitors {default=true}
2018-08-05 16:23:02.505 7fd0fcd23700 10 mds.0.0 create_logger
2018-08-05 16:23:02.505 7fd0fcd23700 7 mds.0.server operator(): full = 0 epoch = 0
2018-08-05 16:23:02.505 7fd0fcd23700 4 mds.0.purge_queue operator(): data pool 3 not found in OS
DMap
2018-08-05 16:23:02.505 7fd0fcd23700 4 mds.0.0 handle_osd_map epoch 0, 0 new blacklist entries
2018-08-05 16:23:02.505 7fd0fcd23700 10 mds.0.server apply_blacklist: killed 0
2018-08-05 16:23:02.505 7fd0fcd23700 10 mds.a handle_mds_map: handling map as rank 0
2018-08-05 16:23:02.505 7fd0fcd23700 1 mds.0.14 handle_mds_map i am now mds.0.14
2018-08-05 16:23:02.505 7fd0fcd23700 1 mds.0.14 handle_mds_map state change up:boot --> up:replay
2018-08-05 16:23:02.505 7fd0fcd23700 10 mds.beacon.a set_want_state: up:standby -> up:replay
2018-08-05 16:23:02.505 7fd0fcd23700 1 mds.0.14 replay_start
2018-08-05 16:23:02.505 7fd0fcd23700 7 mds.0.cache set_recovery_set
2018-08-05 16:23:02.505 7fd0fcd23700 1 mds.0.14 recovery set is
2018-08-05 16:23:02.505 7fd0fcd23700 1 mds.0.14 waiting for osdmap 24 (which blacklists prior in
stance)
2018-08-05 16:23:02.506 7fd0fcd23700 1 -- 172.21.15.26:6817/673291871 <== mon.0 172.21.15.26:6789
/0 28 ==== osd_map(24..24 src has 1..24) v4 ==== 4565+0+0 (79727972 0 0) 0x5635932c2500 con 0x5635
931e4e00
2018-08-05 16:23:02.506 7fd0fcd23700 7 mds.0.server operator(): full = 0 epoch = 24
2018-08-05 16:23:02.506 7fd0fcd23700 4 mds.0.14 handle_osd_map epoch 24, 0 new blacklist entries
2018-08-05 16:23:02.506 7fd0fcd23700 10 mds.0.server apply_blacklist: killed 0
2018-08-05 16:23:02.506 7fd0fcd23700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.26:6789/0 --
mon_subscribe({osdmap=25}) v3 -- 0x56359304f600 con 0
2018-08-05 16:23:02.506 7fd0f6d17700 10 MDSIOContextBase::complete: 12C_IO_Wrapper
2018-08-05 16:23:02.506 7fd0f6d17700 10 MDSInternalContextBase::complete: 15C_MDS_BootStart
2018-08-05 16:23:02.506 7fd0f6d17700 2 mds.0.14 boot_start 0: opening inotable
2018-08-05 16:23:02.506 7fd0f6d17700 10 mds.0.inotable: load
2018-08-05 16:23:02.506 7fd0f6d17700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.187:6804/2215
9 -- osd_op(unknown.0.14:1 2.3 2:c91d4ald::mds0_inotable:head [read 0~0] snapc 0=[] ondisk+read+k
nown_if_redirected+full_force e24) v8 -- 0x563593290840 con 0
2018-08-05 16:23:02.506 7fd0f6d17700 2 mds.0.14 boot_start 0: opening sessionmap
2018-08-05 16:23:02.506 7fd0f6d17700 10 mds.0.sessionmap load
2018-08-05 16:23:02.506 7fd0f6d17700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.187:6804/2215
```

```

9 -- osd_op(unknown.0.14:2 2.3 2:d0630e4c::mds0_sessionmap:head [omap-get-header,omap-get-vals] s
napc 0=[] ondisk+read+known_if_redirected+full_force e24) v8 -- 0x563593290b00 con 0
2018-08-05 16:23:02.507 7fd0f6d17700 2 mds.0.14 boot_start 0: opening mds log
2018-08-05 16:23:02.507 7fd0f6d17700 5 mds.0.log open discovering log bounds
2018-08-05 16:23:02.508 7fd0f6d17700 2 mds.0.14 boot_start 0: opening purge queue (async)
2018-08-05 16:23:02.508 7fd0f5d15700 10 mds.0.log _submit_thread start
2018-08-05 16:23:02.508 7fd0f6d17700 4 mds.0.purge_queue open: opening
2018-08-05 16:23:02.508 7fd0f6d17700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.187:6800/2216
0 -- osd_op(unknown.0.14:3 2.5 2:aa448500::500.00000000:head [read 0~0] snapc 0=[] ondisk+read+kn
own_if_redirected+full_force e24) v8 -- 0x563593290dc0 con 0
2018-08-05 16:23:02.508 7fd0f6d17700 2 mds.0.14 boot_start 0: loading open file table (async)
2018-08-05 16:23:02.508 7fd0f6d17700 10 mds.0.openfiles load
2018-08-05 16:23:02.508 7fd0f6d17700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.26:6813/22088
-- osd_op(unknown.0.14:4 2.6 2:654134d2::mds0_openfiles.0:head [omap-get-header,omap-get-vals] s
napc 0=[] ondisk+read+known_if_redirected+full_force e24) v8 -- 0x563593290840 con 0
2018-08-05 16:23:02.508 7fd0f6d17700 2 mds.0.14 boot_start 0: opening snap table
2018-08-05 16:23:02.508 7fd0f6d17700 10 mds.0.snaptable: load
2018-08-05 16:23:02.508 7fd0f6d17700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.187:6800/2216
0 -- osd_op(unknown.0.14:5 2.5 2:b50e409b::mds_snaptable:head [read 0~0] snapc 0=[] ondisk+read+k
nown_if_redirected+full_force e24) v8 -- 0x563593290b00 con 0
2018-08-05 16:23:02.510 7fd0f6516700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.187:6812/2216
6 -- osd_op(unknown.0.14:6 2.7 2:f1f69726::400.00000000:head [read 0~0] snapc 0=[] ondisk+read+kn
own_if_redirected+full_force e24) v8 -- 0x563593291340 con 0
2018-08-05 16:23:02.517 7fd0ff528700 1 -- 172.21.15.26:6817/673291871 <== osd.7 172.21.15.187:680
0/22160 1 ==== osd_op_reply(3 500.00000000 [read 0~90] v0'0 uv12 ondisk = 0) v8 ==== 156+0+90 (325
0274311 0 1802490552) 0x563593290b00 con 0x5635931e7800
2018-08-05 16:23:02.517 7fd0ff528700 1 -- 172.21.15.26:6817/673291871 <== osd.7 172.21.15.187:680
0/22160 2 ==== osd_op_reply(5 mds_snaptable [read 0~70] v0'0 uv6 ondisk = 0) v8 ==== 157+0+70 (158
2855313 0 294686788) 0x563593290b00 con 0x5635931e7800
2018-08-05 16:23:02.517 7fd0f7d19700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.26:6809/22089
-- osd_op(unknown.0.14:7 2.1 2:86ee9bce::500.00000001:head [stat] snapc 0=[] ondisk+read+rworder
ed+known_if_redirected+full_force e24) v8 -- 0x563593291600 con 0
2018-08-05 16:23:02.517 7fd0f7d19700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.187:6812/2216
6 -- osd_op(unknown.0.14:8 2.7 2:eaf075c4::500.00000002:head [stat] snapc 0=[] ondisk+read+rworde
red+known_if_redirected+full_force e24) v8 -- 0x5635932918c0 con 0
2018-08-05 16:23:02.517 7fd0f6d17700 10 MDSIOContextBase::complete: 12C_IO_MT_Load
2018-08-05 16:23:02.518 7fd0f6d17700 10 mds.0.snaptable: load_2 got 70 bytes
2018-08-05 16:23:02.518 7fd0f6d17700 10 mds.0.snaptable: load_2 loaded v1
2018-08-05 16:23:02.521 7fd0ff528700 1 -- 172.21.15.26:6817/673291871 <== osd.2 172.21.15.26:6809
/22089 1 ==== osd_op_reply(7 500.00000001 [stat] v0'0 uv11 ondisk = 0) v8 ==== 156+0+16 (187129955
3 0 2175828170) 0x563593290580 con 0x5635932efc00
2018-08-05 16:23:02.523 7fd0ffd29700 1 -- 172.21.15.26:6817/673291871 <== osd.6 172.21.15.187:681
2/22166 1 ==== osd_op_reply(6 400.00000000 [read 0~22] v0'0 uv5 ondisk = 0) v8 ==== 156+0+22 (4241
071413 0 3825830296) 0x5635932918c0 con 0x5635932ef500
2018-08-05 16:23:02.523 7fd0ffd29700 1 -- 172.21.15.26:6817/673291871 <== osd.6 172.21.15.187:681
2/22166 2 ==== osd_op_reply(8 500.00000002 [stat] v0'0 uv0 ondisk = -2 ((2) No such file or direct
ory)) v8 ==== 156+0+0 (3018570095 0 0) 0x5635932918c0 con 0x5635932ef500
2018-08-05 16:23:02.523 7fd0f7d19700 4 mds.0.purge_queue operator(): open complete
2018-08-05 16:23:02.524 7fd0f6516700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.26:6813/22088
-- osd_op(unknown.0.14:9 2.4 2:292cf221::200.00000000:head [read 0~0] snapc 0=[] ondisk+read+kno
wn_if_redirected+full_force e24) v8 -- 0x5635932fe000 con 0
2018-08-05 16:23:02.524 7fd0f6516700 4 mds.0.log Waiting for journal 0x200 to recover...
2018-08-05 16:23:02.839 7fd0f9d1d700 10 mds.beacon.a _send up:replay seq 17
2018-08-05 16:23:02.839 7fd0f9d1d700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.26:6789/0 --
mdsbeacon(4560/a up:replay seq 17 v14) v7 -- 0x5635932bca80 con 0
2018-08-05 16:23:02.839 7fd0fcd23700 1 -- 172.21.15.26:6817/673291871 <== mon.0 172.21.15.26:6789
/0 29 ==== mdsbeacon(4560/a up:replay seq 17 v14) v7 ==== 126+0+0 (3387345207 0 0) 0x5635932bca80
con 0x5635931e4e00
2018-08-05 16:23:02.839 7fd0fcd23700 10 mds.beacon.a handle_mds_beacon up:replay seq 17 rtt 0.0005
22
2018-08-05 16:23:04.840 7fd0fad1f700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.26:6800/21275
-- mgrreport(unknown.a +24-0 packed 214) v6 -- 0x5635932b2900 con 0
2018-08-05 16:23:06.097 7fd0ffd29700 1 -- 172.21.15.26:6817/673291871 <== osd.5 172.21.15.187:680
4/22159 1 ==== osd_op_reply(1 mds0_inotable [read 0~34] v0'0 uv2 ondisk = 0) v8 ==== 157+0+34 (209
506393 0 1695280775) 0x5635932918c0 con 0x5635931e7100
2018-08-05 16:23:06.097 7fd0ffd29700 1 -- 172.21.15.26:6817/673291871 <== osd.5 172.21.15.187:680
4/22159 2 ==== osd_op_reply(2 mds0_sessionmap [omap-get-header,omap-get-vals] v0'0 uv3 ondisk = 0)

```

```
v8 ==== 201+0+19 (3119396333 0 1118585302) 0x5635932918c0 con 0x5635931e7100
2018-08-05 16:23:06.097 7fd0f6d17700 10 MDSIOContextBase::complete: 12C_IO_MT_Load
2018-08-05 16:23:06.097 7fd0f6d17700 10 mds.0.inotable: load_2 got 34 bytes
2018-08-05 16:23:06.097 7fd0f6d17700 10 mds.0.inotable: load_2 loaded v0
2018-08-05 16:23:06.097 7fd0f6d17700 10 MDSIOContextBase::complete: N12_GLOBAL__N_112C_IO_SM_LoadE
2018-08-05 16:23:06.097 7fd0f6d17700 10 mds.0.sessionmap _load_finish loaded version 0
2018-08-05 16:23:06.097 7fd0f6d17700 10 mds.0.sessionmap _load_finish: omap load complete
2018-08-05 16:23:06.097 7fd0f6d17700 10 mds.0.sessionmap _load_finish: v 0, 0 sessions
2018-08-05 16:23:06.097 7fd0f6d17700 10 mds.0.sessionmap dump
2018-08-05 16:23:06.839 7fd0f9d1d700 10 mds.beacon.a _send up:replay seq 18
2018-08-05 16:23:06.839 7fd0f9d1d700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.26:6789/0 --
mdsbeacon(4560/a up:replay seq 18 v14) v7 -- 0x5635932bce00 con 0
2018-08-05 16:23:06.839 7fd0fcd23700 1 -- 172.21.15.26:6817/673291871 <== mon.0 172.21.15.26:6789
/0 30 ==== mdsbeacon(4560/a up:replay seq 18 v14) v7 ==== 126+0+0 (3517479427 0 0) 0x5635932bce00
con 0x5635931e4e00
2018-08-05 16:23:06.839 7fd0fcd23700 10 mds.beacon.a handle_mds_beacon up:replay seq 18 rtt 0.0005
68
2018-08-05 16:23:09.840 7fd0fad1f700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.26:6800/21275
-- mgrreport(unknown.a +0-0 packed 214) v6 -- 0x5635932b2c00 con 0
2018-08-05 16:23:10.839 7fd0f9d1d700 10 mds.beacon.a _send up:replay seq 19
2018-08-05 16:23:10.839 7fd0f9d1d700 1 -- 172.21.15.26:6817/673291871 --> 172.21.15.26:6789/0 --
mdsbeacon(4560/a up:replay seq 19 v14) v7 -- 0x5635932bd180 con 0
2018-08-05 16:23:10.840 7fd0fcd23700 1 -- 172.21.15.26:6817/673291871 <== mon.0 172.21.15.26:6789
/0 31 ==== mdsbeacon(4560/a up:replay seq 19 v14) v7 ==== 126+0+0 (3650229487 0 0) 0x5635932bd180
con 0x5635931e4e00
2018-08-05 16:23:10.840 7fd0fcd23700 10 mds.beacon.a handle_mds_beacon up:replay seq 19 rtt 0.0008
99
2018-08-05 16:23:14.839 7fd0f9d1d700 10 mds.beacon.a _send up:replay seq 20
```

From:
/ceph/teuthology-archive/pdonnell-2018-08-04_02:59:35-kcephfs-wip-pdonnell-testing-20180804.001849-testing-basic-smithi/2862358/remote/smithi026/log/ceph-mds.a.log.gz

```
2018-08-05T16:30:02.487 INFO:teuthology.orchestra.run.smithi026.stdout:{"epoch":14,"default_fscid":1,"compat":{"compat":{},"ro_compat":{},"incompat":{"feature_1":"base v0.20","feature_2":"client writeable ranges","feature_3":"default file layouts on dirs","feature_4":"dir inode in separate object","feature_5":"mds uses versioned encoding","feature_6":"dirfrag is stored in omap","feature_8":"no anchor table","feature_9":"file layout v2","feature_10":"snaprealm v2"}},{"feature_flags":{"enable_multiple":false,"ever_enabled_multiple":false},"standbys":[{"gid":4235,"name":"x-s","rank":-1,"incarnation":0,"state":"up:standby","state_seq":1,"addr":"172.21.15.26:6818/4274319287","addrs":{"addrvec":[{"type":"legacy","addr":"172.21.15.26:6818","nonce":4274319287}]},"standby_for_rank":-1,"standby_for_fscid":-1,"standby_for_name":"","standby_replay":false,"export_targets":[],"features":{"4611087854031142911","epoch":5}},"filesystems":[{"mdsmap":{"epoch":14,"flags":18,"ever_allowed_features":0,"explicitly_allowed_features":0,"created":"2018-08-05 16:20:31.431908","modified":"2018-08-05 16:23:02.498495","tableserver":0,"root":0,"session_timeout":60,"session_autoclose":300,"min_compat_client":"-1 (unspecified)","max_file_size":1099511627776,"last_failure":0,"last_failure_osd_epoch":24,"compat":{"compat":{},"ro_compat":{},"incompat":{"feature_1":"base v0.20","feature_2":"client writeable ranges","feature_3":"default file layouts on dirs","feature_4":"dir inode in separate object","feature_5":"mds uses versioned encoding","feature_6":"dirfrag is stored in omap","feature_8":"no anchor table","feature_9":"file layout v2","feature_10":"snaprealm v2"}},{"max_mds":1,"in":[0],"up":{"mds_0":4560},"failed":[],"damaged":[],"stopped":[],"info":{"gid_4560":{"gid":4560,"name":"a","rank":0,"incarnation":14,"state":"up:replay","state_seq":1,"addr":"172.21.15.26:6817/673291871","addrs":{"addrvec":[{"type":"legacy","addr":"172.21.15.26:6817","nonce":673291871}]},"standby_for_rank":-1,"standby_for_fscid":-1,"standby_for_name":"","standby_replay":false,"export_targets":[],"features":{"4611087854031142911}},"data_pools":[3],"metadata_pool":2,"enabled":true,"fs_name":"cephfs","balancer":"","standby_count_wanted":0},"id":1}]}
2018-08-05T16:30:02.488 ERROR:tasks.mds_thrash.fs.[cephfs]:exception:
Traceback (most recent call last):
  File "/home/teuthworker/src/git.ceph.com_ceph-c_wip-pdonnell-testing-20180804.001849/qa/tasks/mds_thrash.py", line 229, in _run
    self.do_thrash()
  File "/home/teuthworker/src/git.ceph.com_ceph-c_wip-pdonnell-testing-20180804.001849/qa/tasks/mds_thrash.py", line 416, in do_thrash
    status = self.wait_for_stable(rank, gid)
```

```
File "/home/teuthworker/src/git.ceph.com_ceph-c_wip-pdonnell-testing-20180804.001849/qa/tasks/mds_thrash.py", line 320, in wait_for_stable
    raise RuntimeError('timeout waiting for cluster to stabilize')
RuntimeError: timeout waiting for cluster to stabilize
2018-08-05T16:30:02.517 INFO:tasks.ceph.osd.5.smithi187.stderr:2018-08-05 16:30:02.581 7fc37c73970
0 -1 osd.5 24 get_health_metrics reporting 498 slow ops, oldest is osd_op(client.4509.0:44485 3.5s
0 3:a6dfaf7c:::1000000171.00000018:head [write 3969024~8192,write 4009984~4096,write 4026368~4096
,write 4038656~8192,write 4055040~4096,write 4067328~4096,write 4083712~4096,write 4112384~4096,wri
te 4124672~4096,write 4165632~4096,write 4173824~8192] snapc l=[] ondisk+write+known_if_redirected
e23)
2018-08-05T16:30:02.611 INFO:tasks.ceph.osd.4.smithi187.stderr:2018-08-05 16:30:02.675 7fee29eee70
0 -1 osd.4 24 get_health_metrics reporting 1342 slow ops, oldest is osd_op(client.4509.0:39257 3.3
s0 3:c3453f39:::100000001a5.00000008:head [write 1409024~4096,write 1425408~4096,write 1437696~409
6,write 1474560~4096,write 1495040~8192,write 1507328~4096,write 1519616~4096,write 1527808~4096,w
rite 1581056~12288,write 1597440~4096,write 1613824~4096,write 1622016~4096,write 1630208~4096,wri
te 1642496~8192,write 1654784~4096,write 1667072~8192] snapc l=[] ondisk+write+known_if_redirected
e23)
2018-08-05T16:30:03.519 INFO:tasks.ceph.osd.5.smithi187.stderr:2018-08-05 16:30:03.584 7fc37c73970
0 -1 osd.5 24 get_health_metrics reporting 496 slow ops, oldest is osd_op(client.4509.0:44503 3.5s
0 3:aae6e7ed:::10000000175.00000001:head [write 225280~4096,write 233472~16384,write 258048~12288,
write 278528~4096,write 286720~8192,write 299008~4096,write 307200~28672,write 344064~4096,write 3
56352~28672,write 389120~20480,write 425984~4096,write 434176~16384,write 454656~8192,write 491520
~8192,write 503808~8192,write 516096~8192] snapc l=[] ondisk+write+known_if_redirected e23)
2018-08-05T16:30:03.616 INFO:tasks.ceph.osd.4.smithi187.stderr:2018-08-05 16:30:03.678 7fee29eee70
0 -1 osd.4 24 get_health_metrics reporting 1342 slow ops, oldest is osd_op(client.4509.0:39257 3.3
s0 3:c3453f39:::100000001a5.00000008:head [write 1409024~4096,write 1425408~4096,write 1437696~409
6,write 1474560~4096,write 1495040~8192,write 1507328~4096,write 1519616~4096,write 1527808~4096,w
rite 1581056~12288,write 1597440~4096,write 1613824~4096,write 1622016~4096,write 1630208~4096,wri
te 1642496~8192,write 1654784~4096,write 1667072~8192] snapc l=[] ondisk+write+known_if_redirected
e23)
2018-08-05T16:30:03.893 INFO:tasks.ceph.osd.3.smithi026.stderr:2018-08-05 16:30:03.445 7fa492ab570
0 -1 osd.3 24 get_health_metrics reporting 6348 slow ops, oldest is osd_op(client.4509.0:21130 3.7
s0 3:fe57b2cc:::10000000174.00000002:head [write 331776~16384,write 356352~16384,write 389120~8192
,write 401408~4096,write 413696~8192,write 425984~4096,write 438272~8192,write 454656~20480,write
483328~8192,write 499712~8192,write 512000~12288,write 532480~12288,write 548864~8192,write 561152
~4096,write 573440~12288,write 589824~4096] snapc l=[] ondisk+write+known_if_redirected e23)
2018-08-05T16:30:03.956 INFO:tasks.ceph.osd.1.smithi026.stderr:2018-08-05 16:30:03.404 7f5aca0a370
0 -1 osd.1 24 get_health_metrics reporting 6968 slow ops, oldest is osd_op(client.4509.0:26739 3.4
s0 3:38f90dac:::100000001c5.00000007:head [write 1363968~8192,write 1380352~4096,write 1400832~409
6,write 1421312~8192,write 1437696~4096,write 1454080~4096,write 1462272~8192,write 1486848~4096,w
rite 1507328~4096,write 1515520~8192,write 1531904~12288,write 1552384~4096,write 1568768~4096,wri
te 1581056~4096,write 1589248~4096,write 1617920~16384] snapc l=[] ondisk+write+known_if_redirected
e23)
2018-08-05T16:30:04.290 INFO:tasks.mds_thrash.daemon_watchdog:thrasher on fs.cephfs failed
2018-08-05T16:30:04.291 INFO:tasks.mds_thrash.daemon_watchdog:BARK! unmounting mounts and killing
all daemons
```

From:
/ceph/teuthology-archive/pdonnell-2018-08-04_02:59:35-kcephfs-wip-pdonnell-testing-20180804.001849-testing-basic-smithi/2862358/teuthology.log

History

#1 - 08/07/2018 09:35 AM - Zheng Yan

```
1148> 2018-08-05 16:23:02.524 7fd06516700 - 172.21.15.26:6817/673291871 --> 172.21.15.26:6813/22088 -- osd_op(unknown.0.14:9 2.4
2:292cf221:::200.00000000:head [read 0~0] snapc 0=] ondisk+read+known_if_redirected+full_force e24) v8 -- 0x5635932fe000 con 0
```

mds didn't receive reply for request that reads journal head

in log of osd.1

```
2018-08-05 16:23:33.650 7f5aca0a3700 20 slow op osd_op(mds.0.14:9 2.4 2.844f3494 (undecoded) ondisk+read+known_if_redirected+full_force
e24) initiated 2018-08-05 16:23:02.527920
```

#2 - 08/13/2018 01:45 PM - Patrick Donnelly

- Assignee changed from Zheng Yan to Patrick Donnelly

#3 - 08/16/2018 09:13 PM - Patrick Donnelly

I think the potential right solution here is to make the thrasher more tolerant of slow recovery if the cluster health shows slow ops. (It's valuable to confirm the MDS can still recover when metadata ops are very slow.) What do you think Zheng?

#4 - 09/12/2018 01:21 PM - Patrick Donnelly

- Priority changed from Urgent to High

#5 - 09/18/2018 06:32 PM - Patrick Donnelly

From Luminous testing:

/ceph/teuthology-archive/yuriw-2018-09-17_23:29:33-kcephfs-wip-yuri2-testing-2018-09-17-1940-luminous-testing-basic-smithi/3035324/remote/smithi017/log/ceph-mds.a.log.gz

#6 - 03/07/2019 11:21 PM - Patrick Donnelly

- Target version changed from v14.0.0 to v15.0.0