

fs - Bug #19406

MDS server crashes due to inconsistent metadata.

03/29/2017 07:57 AM - Christoffer Lilja

Status:	Resolved	Start date:	03/29/2017
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:		Affected Versions:	v10.2.6
Source:		ceph-qa-suite:	fs
Tags:		Component(FS):	MDS, MDSMonitor, ceph-fuse
Backport:	jewel, kraken	Labels (FS):	
Regression:	No	Pull request ID:	
Severity:	3 - minor		
Reviewed:			

Description

```
48> 2017-03-26 23:34:38.886779 7f4b95030700 5 mds.neutron handle_mds_map epoch 620 from mon.1
47> 2017-03-26 23:34:38.886811 7f4b95030700 1 mds.0.618 handle_mds_map i am now mds.0.618
46> 2017-03-26 23:34:38.886817 7f4b95030700 1 mds.0.618 handle_mds_map state change up:reconnect -> up:rejoin
45> 2017-03-26 23:34:38.886829 7f4b95030700 1 mds.0.618 rejoin_start
44> 2017-03-26 23:34:38.889402 7f4b95030700 1 mds.0.618 rejoin_joint_start
43> 2017-03-26 23:34:38.890743 7f4b95030700 1 mds.0.618 rejoin_done
42> 2017-03-26 23:34:38.891686 7f4b95030700 3 mds.0.618 request_state up:active
41> 2017-03-26 23:34:38.891717 7f4b95030700 10 monclient: send_mon_message to mon.orbit at 10.12.0.30:6789/0
40> 2017-03-26 23:34:38.891727 7f4b95030700 1 10.12.0.33:6818/3867961855 --> 10.12.0.30:6789/0 --
mdsbeacon(324154/neutron up:active seq 4 v620) v7 -- ?+0 0x55772168f500 con 0x5577215fa300
39> 2017-03-26 23:34:38.891761 7f4b95030700 5 mds.0.bal rebalance done
38> 2017-03-26 23:34:38.891811 7f4b95030700 1 10.12.0.33:6818/3867961855 <== mon.1 10.12.0.30:6789/0 16 ====
mdsbeacon(324154/neutron up:rejoin seq 3 v620) v7 ==== 132+0+0 (332329670 0 0) 0x55772168f180 con 0x5577215fa300
37> 2017-03-26 23:34:38.891841 7f4b95030700 1 10.12.0.33:6818/3867961855 <== mon.1 10.12.0.30:6789/0 17 ==== log(last
1) v1 ==== 28+0+0 (2312028689 0 0) 0x5577216c1680 con 0x5577215fa300
36> 2017-03-26 23:34:38.891855 7f4b95030700 10 log_client handle_log_ack log(last 1) v1
35> 2017-03-26 23:34:38.891858 7f4b95030700 10 log_client logged 2017-03-26 23:34:36.924844 mds.0
10.12.0.33:6818/3867961855 1 : cluster [ERR] journal replay inotablev mismatch 5 -> 17
34> 2017-03-26 23:34:39.977418 7f4b95030700 1 10.12.0.33:6818/3867961855 <== mon.1 10.12.0.30:6789/0 18 ====
mdsmap(e 621) v1 ==== 691+0+0 (2593754737 0 0) 0x5577216c18c0 con 0x5577215fa300
33> 2017-03-26 23:34:39.977452 7f4b95030700 5 mds.neutron handle_mds_map epoch 621 from mon.1
32> 2017-03-26 23:34:39.977537 7f4b95030700 1 mds.0.618 handle_mds_map i am now mds.0.618
31> 2017-03-26 23:34:39.977549 7f4b95030700 1 mds.0.618 handle_mds_map state change up:rejoin -> up:active
30> 2017-03-26 23:34:39.977557 7f4b95030700 1 mds.0.618 recovery_done successful recovery!
29> 2017-03-26 23:34:39.977561 7f4b95030700 1 10.12.0.33:6818/3867961855 --> mds.0 10.12.0.33:6818/3867961855 --
mds_table_request(snaptable server_ready) v1 -- ?+0 0x5577216c1680
28> 2017-03-26 23:34:39.977643 7f4b95030700 1 10.12.0.33:6818/3867961855 --> 10.12.0.33:6801/4685 --
osd_op(mds.0.618:22 5.c5265ab3 100.00000000 [omap-get-header 0~0,omap-get-vals 0~16,getxattr parent] snapc 0=[])
ack+read+known_if_redirected+full_force e1893) v7 -- ?+0 0x5577215eed80 con 0x5577215fb800
27> 2017-03-26 23:34:39.977675 7f4b95030700 1 mds.0.618 active_start
26> 2017-03-26 23:34:39.977801 7f4b95030700 1 10.12.0.33:6818/3867961855 --> 10.12.0.3:6802/31677 --
osd_op(mds.0.618:23 13.a03303fc 100013135c5.000003ff [stat] snapc 0=[]) ack+read+rworderred+known_if_redirected+full_force
e1893) v7 -- ?+0 0x5577215ecd00 con 0x5577215fbf80
25> 2017-03-26 23:34:39.977854 7f4b95030700 4 mds.0.server handle_client_request client_request(client.244165:37317-getattr
pAsLsXsFs #1 2017-03-26 23:26:06.214233 RETRY=25) v3
24> 2017-03-26 23:34:39.977865 7f4b95030700 5 mds.0.server waiting for root
23> 2017-03-26 23:34:39.977871 7f4b95030700 4 mds.0.server handle_client_request client_request(client.264103:40830-rmdir
#10000000535/bh_meta 2017-03-26 23:11:14.287827 RETRY=39) v3
22> 2017-03-26 23:34:39.977878 7f4b95030700 5 mds.0.server waiting for root
21> 2017-03-26 23:34:39.977882 7f4b95030700 4 mds.0.server handle_client_request client_request(client.264103:40831-getattr
pAsLsXsFs #1 2017-03-26 23:26:06.212140 RETRY=25) v3
20> 2017-03-26 23:34:39.977891 7f4b95030700 5 mds.0.server waiting for root
```

```
-19> 2017-03-26 23:34:39.977904 7f4b95030700 3 mds.0.server.handle_client_session client_session(request_renewcaps seq
19719) v1 from client.244165
-18> 2017-03-26 23:34:39.977918 7f4b95030700 1 10.12.0.33:6818/3867961855 --> 10.12.0.33:0/3931259975 --
client_session(renewcaps seq 19719) v2 -- ?+0 0x5577216c1440 con 0x5577215fa900
-17> 2017-03-26 23:34:39.977946 7f4b95030700 3 mds.0.server.handle_client_session client_session(request_renewcaps seq
19713) v1 from client.264103
-16> 2017-03-26 23:34:39.977958 7f4b95030700 1 10.12.0.33:6818/3867961855 --> 10.12.0.30:0/3952805459 --
client_session(renewcaps seq 19713) v2 -- ?+0 0x557721593440 con 0x5577215fbb00
-15> 2017-03-26 23:34:39.977984 7f4b95030700 1 mds.0.618 cluster recovered.
-14> 2017-03-26 23:34:39.977998 7f4b95030700 5 mds.0.bal rebalance done
-13> 2017-03-26 23:34:39.978004 7f4b95030700 4 mds.0.618 set_osd_epoch_barrier: epoch=1893
-12> 2017-03-26 23:34:39.978021 7f4b95030700 1 10.12.0.33:6818/3867961855 <== mon.1 10.12.0.30:6789/0 19 ====
mdsbeacon(324154/neutron up:active seq 4 v621) v7 ==== 132+0+0 (2124777977 0 0) 0x55772168ff80 con 0x5577215fa300
-11> 2017-03-26 23:34:39.978050 7f4b95030700 1 10.12.0.33:6818/3867961855 <== mds.0 10.12.0.33:6818/3867961855 0 ====
mds_table_request(snaptable server_ready) v1 ==== 0+0+0 (0 0 0) 0x5577216c1680 con 0x5577215fa000
-10> 2017-03-26 23:34:39.978714 7f4b9e81e700 1 10.12.0.33:6818/3867961855 <== osd.2 10.12.0.33:6801/4685 7 ====
osd_op_reply(22 100.00000000 [omap-get-header 0~0,omap-get-vals 0~16,getxattr (30)] v0'0 uv4527 ondisk = 0) v7 ====
216+0+4822 (2504460803 0 789424327) 0x5577215ef400 con 0x5577215fb800
-9> 2017-03-26 23:34:39.978805 7f4b90625700 0 mds.0.cache.dir(100) _fetched_badness: got (but i already had) [inode 604
[...2.head] ~mds0/stray4/ auth v142790 f(v7 m2017-03-15 17:26:48.614897 7878=7227+651) n(v62 re2017-03-15 17:26:48.614897
b215967112272 7870=7222+648) (inest lock dirty) (ifile lock dirty) (iversion lock) | dirtyscattered=2 dirfrag=1 0x5577216ac4f0] mode
16704 mtime 2017-03-15 17:26:48.614897
-8> 2017-03-26 23:34:39.978846 7f4b90625700 1 log_channel(cluster) log [ERR]: loaded dup inode 604 [2.head] v142790 at
~mds0/stray4, but inode 604.head v142790 already exists at ~mds0/stray4
-7> 2017-03-26 23:34:39.978859 7f4b90625700 0 mds.0.cache.dir(100) _fetched_badness: got (but i already had) [inode 602
[...2.head] ~mds0/stray2/ auth v143073 f(v6 m2017-03-17 15:15:19.995740 8994=8215+779) n(v69 re2017-03-17 15:15:19.995740
b930314053756 8983=8209+774) (inest lock dirty) (ifile lock dirty) (iversion lock) | dirtyscattered=2 dirfrag=1 0x5577216ab8a0] mode
16704 mtime 2017-03-17 15:15:19.995740
-6> 2017-03-26 23:34:39.978878 7f4b90625700 1 log_channel(cluster) log [ERR]: loaded dup inode 602 [2.head] v143073 at
~mds0/stray2, but inode 602.head v143073 already exists at ~mds0/stray2
-5> 2017-03-26 23:34:39.978884 7f4b90625700 0 mds.0.cache.dir(100) _fetched_badness: got (but i already had) [inode 601
[...2.head] ~mds0/stray1/ auth v143068 f(v8 m2017-03-15 21:37:11.199303 8113=7114+999) n(v62 re2017-03-15 21:37:11.199303
b595062636506 8089=7101+988) (inest lock dirty) (ifile lock dirty) (iversion lock) | dirtyscattered=2 dirfrag=1 0x5577216ab278] mode
16704 mtime 2017-03-15 21:37:11.199303
-4> 2017-03-26 23:34:39.978907 7f4b90625700 1 log_channel(cluster) log [ERR]: loaded dup inode 601 [2.head] v143068 at
~mds0/stray1, but inode 601.head v143068 already exists at ~mds0/stray1
-3> 2017-03-26 23:34:39.979006 7f4b90e26700 1 10.12.0.33:6818/3867961855 --> 10.12.0.33:6801/4685 -- osd_op(mds.0.618:24
5.6b2cdaff 1.00000000 [omap-get-header 0~0,omap-get-vals 0~16,getxattr parent] snapc 0=[])
ack+read+known_if_redirected+full_force e1893) v7 -- ?+0 0x5577215ef0c0 con 0x5577215fb800
-2> 2017-03-26 23:34:39.979680 7f4b9e81e700 1 10.12.0.33:6818/3867961855 <== osd.2 10.12.0.33:6801/4685 8 ====
osd_op_reply(24 1.00000000 [omap-get-header 0~0,omap-get-vals 0~16,getxattr (30)] v0'0 uv2403 ondisk = 0) v7 ====
214+0+2126 (3403140809 0 1478335496) 0x5577215ed040 con 0x5577215fb800
-1> 2017-03-26 23:34:39.979831 7f4b90e26700 0 mds.0.cache creating system inode with ino:601
0> 2017-03-26 23:34:39.981381 7f4b90e26700 -1 mds/MDCache.cc: In function 'void MDCache::add_inode(CInode*)' thread
7f4b90e26700 time 2017-03-26 23:34:39.979849
mds/MDCache.cc: 265: FAILED assert(inode_map.count(in->vino()) == 0)
```

ceph version 10.2.6 (656b5b63ed7c43bd014bcafd81b001959d5f089f)

- 1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x80) [0x55771694af80]
- 2: (()+0x2e8e75) [0x5577165f7e75]
- 3: (MDCache::create_system_inode(inodeno_t, int)+0x14d) [0x5577165f7fcd]
- 4: (MDCache::populate_mydir()+0x5d2) [0x55771665ab62]
- 5: (MDCache::open_root()+0xde) [0x55771665b05e]
- 6: (MDSInternalContextBase::complete(int)+0x18b) [0x557716794fbb]
- 7: (MDSRank::_advance_queues()+0x6a7) [0x557716546c97]
- 8: (MDSRank::ProgressThread::entry()+0x4a) [0x5577165471ca]
- 9: (()+0x76ba) [0x7f4b9b2f46ba]
- 10: (clone()+0x6d) [0x7f4b997b582d]

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
0/ 5 osd
0/ 5 optracker
0/ 5 objclass
1/ 3 filestore
1/ 3 journal
0/ 5 ms
1/ 5 mon
0/10 monc
1/ 5 paxos
0/ 5 tp
1/ 5 auth
1/ 5 crypto
1/ 1 finisher
1/ 5 heartbeatmap
1/ 5 perfcounter
1/ 5 rgw
1/10 civetweb
1/ 5 javaclient
1/ 5 asok
1/ 1 throttle
0/ 0 refs
1/ 5 xio
1/ 5 compressor
1/ 5 newstore
1/ 5 bluestore
1/ 5 bluefs
1/ 3 bdev
1/ 5 kstore
4/ 5 rocksdb
4/ 5 leveldb
1/ 5 kinetic
1/ 5 fuse
2/ 2 (syslog threshold)
1/ 1 (stderr threshold)
max_recent 10000
max_new 1000
log_file /var/log/ceph/ceph-mds-neutron.log
-- end dump of recent events ---
2017-03-26 23:34:39.983877 7f4b90e26700 -1 ** Caught signal (Aborted) *
in thread 7f4b90e26700 thread_name:mds_rank_progr

ceph version 10.2.6 (656b5b63ed7c43bd014bcafd81b001959d5f089f)
1: (()+0x53373e) [0x55771684273e]
2: (()+0x11390) [0x7f4b9b2fe390]
3: (gsignal()+0x38) [0x7f4b996e4428]
4: (abort()+0x16a) [0x7f4b996e602a]
5: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x26b) [0x55771694b16b]
6: (()+0x2e8e75) [0x5577165f7e75]
7: (MDCache::create_system_inode(inodeno_t, int)+0x14d) [0x5577165f7fcd]
8: (MDCache::populate_mydir()+0x5d2) [0x55771665ab62]
9: (MDCache::open_root()+0xde) [0x55771665b05e]

```

```
10: (MDSInternalContextBase::complete(int)+0x18b) [0x557716794fbb]
11: (MDSRank::_advance_queues()+0x6a7) [0x557716546c97]
12: (MDSRank::ProgressThread::entry()+0x4a) [0x5577165471ca]
13: (()+0x76ba) [0x7f4b9b2f46ba]
14: (clone()+0x6d) [0x7f4b997b582d]
NOTE: a copy of the executable, or `objdump -rds &lt;executable&gt;` is needed to interpret this.
```

--- begin dump of recent events ---

```
O> 2017-03-26 23:34:39.983877 7f4b90e26700 -1 ** Caught signal (Aborted) *
in thread 7f4b90e26700 thread_name:mds_rank_progr
```

```
ceph version 10.2.6 (656b5b63ed7c43bd014bcafd81b001959d5f089f)
```

```
1: (()+0x53373e) [0x55771684273e]
2: (()+0x11390) [0x7f4b9b2fe390]
3: (gsignal()+0x38) [0x7f4b996e4428]
4: (abort()+0x16a) [0x7f4b996e602a]
5: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x26b) [0x55771694b16b]
6: (()+0x2e8e75) [0x5577165f7e75]
7: (MDCache::create_system_inode(inodeno_t, int)+0x14d) [0x5577165f7fcd]
8: (MDCache::populate_mydir()+0x5d2) [0x55771665ab62]
9: (MDCache::open_root()+0xde) [0x55771665b05e]
10: (MDSInternalContextBase::complete(int)+0x18b) [0x557716794fbb]
11: (MDSRank::_advance_queues()+0x6a7) [0x557716546c97]
12: (MDSRank::ProgressThread::entry()+0x4a) [0x5577165471ca]
13: (()+0x76ba) [0x7f4b9b2f46ba]
14: (clone()+0x6d) [0x7f4b997b582d]
NOTE: a copy of the executable, or `objdump -rds &lt;executable&gt;` is needed to interpret this.
```

--- logging levels ---

```
O/ 5 none
O/ 1 lockdep
O/ 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
O/ 1 buffer
O/ 1 timer
O/ 1 filer
O/ 1 striper
O/ 1 objecter
O/ 5 rados
O/ 5 rbd
O/ 5 rbd_mirror
O/ 5 rbd_replay
O/ 5 journaler
O/ 5 objectcacher
O/ 5 client
O/ 5 osd
O/ 5 optracker
O/ 5 objclass
1/ 3 filestore
1/ 3 journal
O/ 5 ms
1/ 5 mon
O/10 monc
1/ 5 paxos
O/ 5 tp
1/ 5 auth
1/ 5 crypto
```

```
1/ 1 finisher
1/ 5 heartbeatmap
1/ 5 perfcounter
1/ 5 rgw
1/10 civetweb
1/ 5 javaclient
1/ 5 asok
1/ 1 throttle
0/ 0 refs
1/ 5 xio
1/ 5 compressor
1/ 5 newstore
1/ 5 bluestore
1/ 5 bluefs
1/ 3 bdev
1/ 5 kstore
4/ 5 rocksdb
4/ 5 leveldb
1/ 5 kinetic
1/ 5 fuse
2/ 2 (syslog threshold)
-1/-1 (stderr threshold)
max_recent 10000
max_new 1000
log_file /var/log/ceph/ceph-mds.neutron.log
-- end dump of recent events ---
```

Related issues:

- Copied to fs - Backport #19619: jewel: MDS server crashes due to inconsistent... **Resolved**
- Copied to fs - Backport #19620: kraken: MDS server crashes due to inconsisten... **Resolved**

History

#1 - 03/29/2017 09:43 AM - Zheng Yan

please set 'debug_mds = 20', restart the mds and upload the mds.log

#2 - 03/29/2017 10:15 AM - John Spray

Please also tell us what else was going on around this event. Had you recently upgraded? Had another daemon recently failed? Had you run any special commands?

#3 - 03/29/2017 10:44 AM - Christoffer Lilja

- File *ceph-mds.log_trunc* added

A NIC broke and cause tremendous load on the server. Shortly after I had to kill the server by power, note that this is a very old server that easily get overloaded.

I got the error "mds0: metadata damaged" after I got the Ceph cluster online again. It was some dir_frag errors. After that I followed <http://docs.ceph.com/docs/jewel/cephfs/disaster-recovery> to see if that would fix the error.

The executed steps did not solve the errors so I deleted all data I could (Some folders got the "Dir not empty") and started copy in all data again to the cephfs-fuse mount using rsync. After this the mds on the very same node crashes like this.

#4 - 03/29/2017 10:55 AM - Christoffer Lilja

I see now that the two other MDS servers crash the same way.

#5 - 03/29/2017 12:43 PM - John Spray

Can you go back to your logs from when the original metadata damage occurred? There should be a message in your cluster log, and perhaps some more detailed information in the MDS log from the same timestamp.

The disaster recovery tools can make things worse as well as better -- in this instance that seems to be the case.

#6 - 03/29/2017 01:01 PM - John Spray

To work out how the metadata is now corrupted (i.e. what the recovery tools broke), it would be useful to have the full log from MDS startup through to the point where it crashes. In that truncated snippet I can see that it is failing because an inode is getting loaded twice, but I can't see the first time it was loaded.

If the amount of metadata in your pool is not too large, and it isn't confidential, you could also do a "rados -p <metadata pool> export <file>" and send me that.

#7 - 03/29/2017 01:28 PM - Christoffer Lilja

The logfile contains file structure that in it self is sensitive information, that's why it's truncated. I'll see if I can trace the inode that is loaded twice.

#8 - 03/29/2017 02:02 PM - Zheng Yan

looks like ~mds0's dir_layout.dl_dir_hash is wrong. Please check if below patch can prevent mds from crash

```
diff --git a/src/mds/CInode.cc b/src/mds/CInode.cc
index 836362eaad..0bf6b18954 100644
--- a/src/mds/CInode.cc
+++ b/src/mds/CInode.cc
@@ -498,6 +498,10 @@ __u32 InodeStoreBase::hash_dentry_name(const string &dn)
     int which = inode.dir_layout.dl_dir_hash;
     if (!which)
         which = CEPH_STR_HASH_LINUX;
+   if (which != CEPH_STR_HASH_RJENKINS) {
+       dout(0) << " bad dir_layout.dl_dir_hash " << dendl;
+       which = CEPH_STR_HASH_RJENKINS;
+   }
     assert(ceph_str_hash_valid(which));
     return ceph_str_hash(which, dn.data(), dn.length());
 }
```

#9 - 03/29/2017 06:45 PM - John Spray

Looking at what DataScan.cc currently does, looks like we were injecting things with hash set to zero, so it would pick up HASH_LINUX when loaded by the MDS. In case that's what's going on here, I guess we should change the tool to avoid doing that in future:

<https://github.com/ceph/ceph/pull/14234>

Zheng: can you explain what seems to be happening here? It's not obvious to me, as the ~mds0 and stray dirs should usually not be fragmented, so

the hash function shouldn't make any difference to which object it's trying to look up dentries in?

#10 - 03/29/2017 07:09 PM - Christoffer Lilja

I don't have a build environment nor the Ceph source code on the machines.
I'll download the source code, apply the patch and try to compile it in the upcoming days.

And I'm sorry to say that the logs from the event that was the culmination to this has been rotated and deleted.

#11 - 03/30/2017 01:23 AM - Zheng Yan

```
-78> 2017-03-29 12:11:06.298814 7f5e16746700 20 mds.0.cache.dir(100) _fetched pos 4 marker 'I' dname 'stray
4 [2,head]
-77> 2017-03-29 12:11:06.298815 7f5e16746700 20 mds.0.cache.dir(100) lookup (head, 'stray4')
-76> 2017-03-29 12:11:06.298819 7f5e16746700 0 mds.0.cache.dir(100) _fetched badness: got (but i already
had) [inode 604 [...2,head] ~mds0/stray4/ auth v142790 f(v7 m2017-03-15 17:26:48.614897 7878=7227+651) n(v62 r
c2017-03-15 17:26:48.614897 b215967112272 7870=7222+648) (inest lock dirty) (ifile lock dirty) (iversion lock)
| dirtyscattered=2 dirfrag=1 0x55cb6ba3e4f0] mode 16704 mtime 2017-03-15 17:26:48.614897
-75> 2017-03-29 12:11:06.298850 7f5e16746700 -1 log_channel(cluster) log [ERR] : loaded dup inode 604 [2,he
ad] v142790 at ~mds0/stray4, but inode 604.head v142790 already exists at ~mds0/stray4
```

there is a ~mds0/stray4 dentry, but lookup can't find it. The can happen if the hash type for inserting the dentry is different from the hash type for lookup.

#12 - 04/01/2017 06:15 PM - Christoffer Lilja

I downloaded the 10.2.6 tar and tried to apply the above patch from Zheng Yan but it got rejected. I guess it's written for another version than I use?

#13 - 04/05/2017 02:03 AM - Zheng Yan

Christoffer Lilja wrote:

I downloaded the 10.2.6 tar and tried to apply the above patch from Zheng Yan but it got rejected. I guess it's written for another version than I use?

I wrote the patch for the newest code. The patch just adds 3 lines, you can easily edit the src/mds/CInode.cc manually.

#14 - 04/05/2017 07:54 PM - Christoffer Lilja

Hi,

I have little spare time to spend on this and can't get past this compile error due to my very limited knowledge in C++.

May you help me out figure out what's the error in "return ((int*)&__s64)[argc];"?

Run Build Command: "/usr/bin/make" "cmTC_d0692/fast"

/usr/bin/make -f CMakeFiles/cmTC_d0692.dir/build.make CMakeFiles/cmTC_d0692.dir/build

make¹: Entering directory '/home/zidz/workspace/ceph/ceph-10.2.6/build/CMakeFiles/CMakeTmp'

Building C object CMakeFiles/cmTC_d0692.dir/CheckSymbolExists.c.o

/usr/bin/cc -o CMakeFiles/cmTC_d0692.dir/CheckSymbolExists.c.o -c

/home/zidz/workspace/ceph/ceph-10.2.6/build/CMakeFiles/CMakeTmp/CheckSymbolExists.c

/home/zidz/workspace/ceph/ceph-10.2.6/build/CMakeFiles/CMakeTmp/CheckSymbolExists.c: In function 'main':

/home/zidz/workspace/ceph/ceph-10.2.6/build/CMakeFiles/CMakeTmp/CheckSymbolExists.c:9:19: error: expected expression before '__s64'

return ((int*)&__s64)[argc];

^

CMakeFiles/cmTC_d0692.dir/build.make:65: recipe for target 'CMakeFiles/cmTC_d0692.dir/CheckSymbolExists.c.o' failed

make¹: * [CMakeFiles/cmTC_d0692.dir/CheckSymbolExists.c.o] Error 1

make¹: Leaving directory '/home/zidz/workspace/ceph/ceph-10.2.6/build/CMakeFiles/CMakeTmp'

Makefile:126: recipe for target 'cmTC_d0692/fast' failed

make: * [cmTC_d0692/fast] Error 2

File /home/zidz/workspace/ceph/ceph-10.2.6/build/CMakeFiles/CMakeTmp/CheckSymbolExists.c:

/* */

#include <sys/types.h>

#include <linux/types.h>

int main(int argc, char** argv) {

(void)argv;

#ifndef __s64

return ((int*)&__s64)[argc];

#else

(void)argv;

return 0;

#endif

}

#15 - 04/06/2017 01:03 AM - Zheng Yan

Christoffer Lilja wrote:

Hi,

I have little spare time to spend on this and can't get past this compile error due to my very limited knowledge in C++.

May you help me out figure out what's the error in "return ((int*)&__s64)[argc];"?

```
Run Build Command:"/usr/bin/make" "cmTC_d0692/fast"
/usr/bin/make -f CMakeFiles/cmTC_d0692.dir/build.make CMakeFiles/cmTC_d0692.dir/build
make1: Entering directory '/home/zidz/workspace/ceph/ceph-10.2.6/build/CMakeFiles/CMakeTmp'
Building C object CMakeFiles/cmTC_d0692.dir/CheckSymbolExists.c.o
/usr/bin/cc -o CMakeFiles/cmTC_d0692.dir/CheckSymbolExists.c.o -c
/home/zidz/workspace/ceph/ceph-10.2.6/build/CMakeFiles/CMakeTmp/CheckSymbolExists.c
/home/zidz/workspace/ceph/ceph-10.2.6/build/CMakeFiles/CMakeTmp/CheckSymbolExists.c: In function 'main':
/home/zidz/workspace/ceph/ceph-10.2.6/build/CMakeFiles/CMakeTmp/CheckSymbolExists.c:9:19: error: expected expression before '__s64'
return ((int*)&__s64)[argc];
      ^
CMakeFiles/cmTC_d0692.dir/build.make:65: recipe for target 'CMakeFiles/cmTC_d0692.dir/CheckSymbolExists.c.o' failed
make1: * [CMakeFiles/cmTC_d0692.dir/CheckSymbolExists.c.o] Error 1
make1: Leaving directory '/home/zidz/workspace/ceph/ceph-10.2.6/build/CMakeFiles/CMakeTmp'
Makefile:126: recipe for target 'cmTC_d0692/fast' failed
make: * [cmTC_d0692/fast] Error 2
```

```
File /home/zidz/workspace/ceph/ceph-10.2.6/build/CMakeFiles/CMakeTmp/CheckSymbolExists.c:
/* */
#include <sys/types.h>
#include <linux/types.h>

int main(int argc, char** argv) {
(void)argv;
#ifdef __s64
return ((int*)&__s64)[argc];
#else
(void)argv;
return 0;
#endif
}
```

your machine misses some packages that are required for build ceph. try run ./install_deps.sh in ceph source first

#16 - 04/08/2017 09:41 AM - Christoffer Lilja

Okay, I followed the guide at the Ceph site how to prepare and compile Ceph.
I already did run `install_deps.sh` in order do have all packages installed when I'd start compile.
Maybe it's missing packages for Ubuntu 16.04.

#17 - 04/08/2017 09:52 AM - Christoffer Lilja

I'll download latest master instead and try.

#18 - 04/08/2017 10:03 AM - Christoffer Lilja

- *File log.txt added*

- *File CMakeError.log added*

The master `install-deps.sh` installed a flood of other stuff the earlier didn't. But still get fatal errors when running `cmake`, see `log.txt` and `CMakeError.log`

#19 - 04/13/2017 11:07 AM - John Spray

- *Status changed from New to Pending Backport*

Marking as pending backport for the fix to `data-scan` which seems likely to be the underlying cause <https://github.com/ceph/ceph/pull/14234>

#20 - 04/13/2017 11:07 AM - John Spray

- *Backport set to jewel, kraken*

#21 - 04/13/2017 02:31 PM - Nathan Cutler

- *Copied to Backport #19619: jewel: MDS server crashes due to inconsistent metadata. added*

#22 - 04/13/2017 02:31 PM - Nathan Cutler

- *Copied to Backport #19620: kraken: MDS server crashes due to inconsistent metadata. added*

#23 - 04/14/2017 09:59 PM - Nathan Cutler

master PR: <https://github.com/ceph/ceph/pull/14234>

#24 - 09/05/2017 03:37 PM - Nathan Cutler

- *Status changed from Pending Backport to Resolved*

#25 - 09/18/2017 09:53 AM - demiao wu

Christoffer Lilja wrote:

The master `install-deps.sh` installed a flood of other stuff the earlier didn't. But still get fatal errors when running `cmake`, see `log.txt` and `CMakeError.log`

how to solve the build error? I also met this error. thank

Files

ceph-mds.log_trunc	42.1 KB	03/29/2017	Christoffer Lilja
log.txt	22.2 KB	04/08/2017	Christoffer Lilja
CMakeError.log	7.24 KB	04/08/2017	Christoffer Lilja