

Ceph - Bug #13742

MonitorDBStore::get semantic change triggers regressions

11/10/2015 07:27 AM - Loic Dachary

Status:	Resolved	Start date:	11/10/2015
Priority:	Urgent	Due date:	
Assignee:	Joao Eduardo Luis	% Done:	0%
Category:	Monitor	Estimated time:	0.00 hour
Target version:		Spent time:	0.00 hour
Source:	other	Reviewed:	
Tags:		Affected Versions:	
Backport:		ceph-qa-suite:	
Regression:	No	Pull request ID:	
Severity:	3 - minor		

Description

On master make check

```
FAIL: test/mon/mkfs.sh
```

```
=====
```

```
+ PS4='${BASH_SOURCE[0]}:$LINENO: ${FUNCNAME[0]}: '
./test/mon/mkfs.sh:21: : DIR=mkfs
./test/mon/mkfs.sh:22: : export CEPH_CONF=/dev/null
./test/mon/mkfs.sh:22: : CEPH_CONF=/dev/null
./test/mon/mkfs.sh:23: : unset CEPH_ARGS
./test/mon/mkfs.sh:24: : MON_ID=a
./test/mon/mkfs.sh:25: : MON_DIR=mkfs/a
./test/mon/mkfs.sh:26: : CEPH_MON=127.0.0.1:7110
./test/mon/mkfs.sh:27: : TIMEOUT=360
./test/mon/mkfs.sh:185: : run
./test/mon/mkfs.sh:172: run: local actions
./test/mon/mkfs.sh:173: run: actions+='makedir '
./test/mon/mkfs.sh:174: run: actions+='idempotent '
./test/mon/mkfs.sh:175: run: actions+='auth_cephx_key '
./test/mon/mkfs.sh:176: run: actions+='auth_cephx_keyring '
./test/mon/mkfs.sh:177: run: actions+='auth_none '
./test/mon/mkfs.sh:178: run: for action in '$actions'
./test/mon/mkfs.sh:179: run: setup
./test/mon/mkfs.sh:30: setup: teardown
./test/mon/mkfs.sh:35: teardown: kill_daemons
../test/mon/mkfs.sh:69: kill_daemons: find mkfs -name pidfile
find: 'mkfs': No such file or directory
./test/mon/mkfs.sh:36: teardown: rm -fr mkfs
./test/mon/mkfs.sh:31: setup: mkdir mkfs
./test/mon/mkfs.sh:180: run: makedir
./test/mon/mkfs.sh:146: makedir: local toodeep=mkfs/a/toodeep
./test/mon/mkfs.sh:150: makedir: ./ceph-mon --id a --mon-osd-full-ratio=.99 --mon-data-avail-crit
=1 --erasure-code-dir=.libs --mkfs --mon-data=mkfs/a/toodeep
./test/mon/mkfs.sh:155: makedir: tee mkfs/makedir.log
mkdir(mkfs/a/toodeep) : (2) No such file or directory
./test/mon/mkfs.sh:156: makedir: grep 'toodeep.*No such file' mkfs/makedir.log
./test/mon/mkfs.sh:157: makedir: rm mkfs/makedir.log
./test/mon/mkfs.sh:160: makedir: mkdir mkfs/a
./test/mon/mkfs.sh:161: makedir: mon_mkfs --auth-supported=none
./test/mon/mkfs.sh:161: makedir: tee mkfs/makedir.log
../test/mon/mkfs.sh:40: mon_mkfs: uuidgen
./test/mon/mkfs.sh:40: mon_mkfs: local fsid=b90842b3-72aa-4def-81e7-873358743ea6
./test/mon/mkfs.sh:43: mon_mkfs: ./ceph-mon --id a --fsid b90842b3-72aa-4def-81e7-873358743ea6 --
```

```

erasure-code-dir=.libs --mkfs --mon-data=mkfs/a --mon-initial-members=a --mon-host=127.0.0.1:7110
--auth-supported=none
./ceph-mon: mon.noname-a 127.0.0.1:7110/0 is local, renaming to mon.a
./ceph-mon: set fsid to b90842b3-72aa-4def-81e7-873358743ea6
./ceph-mon: created monfs at mkfs/a for mon.a
./test/mon/mkfs.sh:162: mkdir: grep 'mkfs/a already exists' mkfs/mkdir.log
./test/mon/mkfs.sh:181: run: teardown
./test/mon/mkfs.sh:35: teardown: kill_daemons
../test/mon/mkfs.sh:69: kill_daemons: find mkfs -name pidfile
./test/mon/mkfs.sh:36: teardown: rm -fr mkfs
./test/mon/mkfs.sh:178: run: for action in '$actions'
./test/mon/mkfs.sh:179: run: setup
./test/mon/mkfs.sh:30: setup: teardown
./test/mon/mkfs.sh:35: teardown: kill_daemons
../test/mon/mkfs.sh:69: kill_daemons: find mkfs -name pidfile
find: 'mkfs': No such file or directory
./test/mon/mkfs.sh:36: teardown: rm -fr mkfs
./test/mon/mkfs.sh:31: setup: mkdir mkfs
./test/mon/mkfs.sh:180: run: idempotent
./test/mon/mkfs.sh:166: idempotent: mon_mkfs --auth-supported=none
../test/mon/mkfs.sh:40: mon_mkfs: uuidgen
./test/mon/mkfs.sh:40: mon_mkfs: local fsid=78631327-eec6-4154-930f-3161cb1c2257
./test/mon/mkfs.sh:43: mon_mkfs: ./ceph-mon --id a --fsid 78631327-eec6-4154-930f-3161cb1c2257 --
erasure-code-dir=.libs --mkfs --mon-data=mkfs/a --mon-initial-members=a --mon-host=127.0.0.1:7110
--auth-supported=none
./ceph-mon: mon.noname-a 127.0.0.1:7110/0 is local, renaming to mon.a
./ceph-mon: set fsid to 78631327-eec6-4154-930f-3161cb1c2257
./ceph-mon: created monfs at mkfs/a for mon.a
./test/mon/mkfs.sh:167: idempotent: mon_mkfs --auth-supported=none
./test/mon/mkfs.sh:167: idempotent: tee mkfs/mkdir.log
../test/mon/mkfs.sh:40: mon_mkfs: uuidgen
./test/mon/mkfs.sh:40: mon_mkfs: local fsid=0b9ea42f-c91e-4a41-a4fb-c49ac26ce0a9
./test/mon/mkfs.sh:43: mon_mkfs: ./ceph-mon --id a --fsid 0b9ea42f-c91e-4a41-a4fb-c49ac26ce0a9 --
erasure-code-dir=.libs --mkfs --mon-data=mkfs/a --mon-initial-members=a --mon-host=127.0.0.1:7110
--auth-supported=none
'mkfs/a' already exists and is not empty: monitor may already exist
./test/mon/mkfs.sh:168: idempotent: grep '\''mkfs/a\'' already exists' mkfs/mkdir.log
./test/mon/mkfs.sh:181: run: teardown
./test/mon/mkfs.sh:35: teardown: kill_daemons
../test/mon/mkfs.sh:69: kill_daemons: find mkfs -name pidfile
./test/mon/mkfs.sh:36: teardown: rm -fr mkfs
./test/mon/mkfs.sh:178: run: for action in '$actions'
./test/mon/mkfs.sh:179: run: setup
./test/mon/mkfs.sh:30: setup: teardown
./test/mon/mkfs.sh:35: teardown: kill_daemons
../test/mon/mkfs.sh:69: kill_daemons: find mkfs -name pidfile
find: 'mkfs': No such file or directory
./test/mon/mkfs.sh:36: teardown: rm -fr mkfs
./test/mon/mkfs.sh:31: setup: mkdir mkfs
./test/mon/mkfs.sh:180: run: auth_cephx_key
./test/mon/mkfs.sh:119: auth_cephx_key: '[' -f /etc/ceph/keyring ']'
../test/mon/mkfs.sh:124: auth_cephx_key: ./ceph-authtool --gen-print-key
./test/mon/mkfs.sh:124: auth_cephx_key: local key=AQBtk0FWWpm9JBAAL43XvZFJ+MqyUxBsv10EKw==
./test/mon/mkfs.sh:126: auth_cephx_key: mon_mkfs '--key=corrupted key'
../test/mon/mkfs.sh:40: mon_mkfs: uuidgen
./test/mon/mkfs.sh:40: mon_mkfs: local fsid=e6adcf34-7aa3-4ed5-82e4-a8acfed61e28
./test/mon/mkfs.sh:43: mon_mkfs: ./ceph-mon --id a --fsid e6adcf34-7aa3-4ed5-82e4-a8acfed61e28 --
erasure-code-dir=.libs --mkfs --mon-data=mkfs/a --mon-initial-members=a --mon-host=127.0.0.1:7110
'--key=corrupted key'
./ceph-mon: mon.noname-a 127.0.0.1:7110/0 is local, renaming to mon.a
./ceph-mon: set fsid to e6adcf34-7aa3-4ed5-82e4-a8acfed61e28
2015-11-10 06:49:17.682898 7fd2b544a780 -1 mon.a@-1(probing) e0 unable to find a keyring file on /
etc/ceph/ceph.mon.a.keyring,/etc/ceph/ceph.keyring,/etc/ceph/keyring,/etc/ceph/keyring.bin: (2) No
such file or directory
2015-11-10 06:49:17.683090 7fd2b544a780 -1 mon.a@-1(probing) e0 error decoding keyring [mon.]
key = corrupted key
caps mon = "allow *"

```

```

: buffer::malformed_input: error setting modifier for [mon.] type=key val=corrupted key
./ceph-mon: error creating monfs: (22) Invalid argument
./test/mon/mkfs.sh:129: auth_cephx_key: rm -fr mkfs/a/store.db
./test/mon/mkfs.sh:132: auth_cephx_key: mon_mkfs --key=AQBtk0FWWpm9JBAAL43XvZFJ+MqyUxBsv10EKw==
../test/mon/mkfs.sh:40: mon_mkfs: uuidgen
./test/mon/mkfs.sh:40: mon_mkfs: local fsid=2ccb2781-00ad-4301-8771-38258a1dec85
./test/mon/mkfs.sh:43: mon_mkfs: ./ceph-mon --id a --fsid 2ccb2781-00ad-4301-8771-38258a1dec85 --
erasure-code-dir=.libs --mkfs --mon-data=mkfs/a --mon-initial-members=a --mon-host=127.0.0.1:7110
--key=AQBtk0FWWpm9JBAAL43XvZFJ+MqyUxBsv10EKw==
./ceph-mon: mon.noname-a 127.0.0.1:7110/0 is local, renaming to mon.a
./ceph-mon: set fsid to 2ccb2781-00ad-4301-8771-38258a1dec85
2015-11-10 06:49:17.742717 7f4c4ef96780 -1 mon.a@-1(probing) e0 unable to find a keyring file on /
etc/ceph/ceph.mon.a.keyring,/etc/ceph/ceph.keyring,/etc/ceph/keyring,/etc/ceph/keyring.bin: (2) No
such file or directory
./ceph-mon: created monfs at mkfs/a for mon.a
./test/mon/mkfs.sh:134: auth_cephx_key: '[' -f mkfs/a/keyring ']'
./test/mon/mkfs.sh:135: auth_cephx_key: grep AQBtk0FWWpm9JBAAL43XvZFJ+MqyUxBsv10EKw== mkfs/a/keyr
ing
key = AQBtk0FWWpm9JBAAL43XvZFJ+MqyUxBsv10EKw==
./test/mon/mkfs.sh:137: auth_cephx_key: mon_run
./test/mon/mkfs.sh:55: mon_run: ./ceph-mon --id a --chdir= --mon-osd-full-ratio=.99 --mon-data-av
ail-crit=1 --erasure-code-dir=.libs --mon-data=mkfs/a --log-file=mkfs/a/log --mon-cluster-log-file
=mkfs/a/log --run-dir=mkfs/a --pid-file=mkfs/a/pidfile --public-addr 127.0.0.1:7110
./test/mon/mkfs.sh:139: auth_cephx_key: timeout 360 ./ceph --name mon. --keyring mkfs/a/keyring -
mon-host 127.0.0.1:7110 mon stat
*** DEVELOPER MODE: setting PATH, PYTHONPATH and LD_LIBRARY_PATH ***
2015-11-10 06:49:17.997410 7fd9bc50e700 0 -- :/2887895518 >> 127.0.0.1:7110/0 pipe(0x7fd9b8063550
sd=3 :0 s=1 pgs=0 cs=0 l=1
...
2015-11-10 06:54:17.997355 7fd9c0488700 0 monclient(hunting): authenticate timed out after 300
2015-11-10 06:54:17.997383 7fd9c0488700 0 librados: mon. authentication error (110) Connection ti
med out
Error connecting to cluster: TimedOut
./test/mon/mkfs.sh:142: auth_cephx_key: return 1
./test/mon/mkfs.sh:180: run: return 1

[0;31m=====
[0;31mTestsuite summary for ceph 9.2.0[m
[0;31m=====

```

History

#1 - 11/10/2015 07:44 AM - Loic Dachary

<https://github.com/ceph/ceph/pull/6312> may have introduced a regression

#2 - 11/10/2015 07:59 AM - Loic Dachary

```
./ceph-mon --id a --chdir= --mon-osd-full-ratio=.99 --mon-data-avail-crit=1 --erasure-code-dir=.libs --mon-data=mkfs/a --log-file=mkfs/a/log
--mon-cluster-log-file=mkfs/a/log --run-dir=mkfs/a --pid-file=mkfs/a/pidfile --public-addr 127.0.0.1:7110
```

```
2015-11-10 08:55:19.362228 7fb8fc118780 -1 mon/AuthMonitor.cc: In function 'virtual void AuthMonitor::create_i
nitial()' thread 7fb8fc118780 time 2015-11-10 08:55:19.361259
mon/AuthMonitor.cc: 101: FAILED assert(ret == 0)
```

```

ceph version 9.2.0-523-g866d0a0 (866d0a0892dcfcb3d636a6fab0da4637e8efb4d1)
1: (ceph::__ceph_assert_fail(char const*, char const*, int, char const*)+0x8b) [0x7fb8fbcd955b]
2: (AuthMonitor::create_initial()+0x53b) [0x7fb8fb44fcb]
3: (PaxosService::_active()+0x516) [0x7fb8fbaebc26]
4: (PaxosService::election_finished()+0x67) [0x7fb8fbaec007]
5: (Monitor::win_election(unsigned int, std::set<int, std::less<int>, std::allocator<int> >&, unsigned long,
MonCommand const*, int, std::set<int, std::less<int>, std::allocator<int> > const*)+0x236) [0x7fb8fbaa6a76]
6: (Monitor::win_standalone_election()+0x14d) [0x7fb8fbaa6e8d]
7: (Monitor::bootstrap()+0xa2b) [0x7fb8fbaa79eb]
8: (Monitor::init()+0x184) [0x7fb8fbaa7ce4]
9: (main()+0x23fc) [0x7fb8fba67e8c]
10: (__libc_start_main()+0xf5) [0x7fb8f958cec5]
11: ((()+0x1a309a) [0x7fb8fba7709a]
NOTE: a copy of the executable, or `objdump -rds <executable>` is needed to interpret this.

```

#3 - 11/10/2015 08:11 AM - Loic Dachary

Reverting <https://github.com/ceph/ceph/commit/66b7b920cf5a0a9c71212573ef47fb2c7ea9b5ff>, the test passes

#4 - 11/10/2015 08:31 AM - Loic Dachary

- *Subject changed from test/mon/mkfs.sh fails to MonitorDBStore::get semantic change triggers regressions*

#5 - 11/10/2015 08:33 AM - Loic Dachary

- *Assignee changed from Loic Dachary to Piotr Dalek*

#6 - 11/10/2015 08:34 AM - Loic Dachary

<https://github.com/ceph/ceph/commit/66b7b920cf5a0a9c71212573ef47fb2c7ea9b5ff> changed the behavior of the function. It previously never returned ENOENT because the key was inserted, always and the size of out was always 1. The callers have to be changed to deal with the change.

#7 - 11/10/2015 08:55 AM - Piotr Dalek

Loic Dachary wrote:

<https://github.com/ceph/ceph/commit/66b7b920cf5a0a9c71212573ef47fb2c7ea9b5ff> changed the behavior of the function. It previously never returned ENOENT because the key was inserted, always and the size of out was always 1. The callers have to be changed to deal with the change.

This calls for a larger patch which cannot be delivered fast. Instead, I'll restore the old behavior and then propose a patch that will fix this globally.

#8 - 11/10/2015 11:53 AM - Piotr Dalek

Pushed <https://github.com/ceph/ceph/pull/6522> with temporary fix.

#9 - 11/12/2015 04:09 PM - Loic Dachary

- *Status changed from In Progress to Resolved*

#10 - 11/12/2015 04:13 PM - Loic Dachary

- *Status changed from Resolved to In Progress*

#11 - 11/13/2015 12:09 PM - Piotr Dalek

Pushed <https://github.com/ceph/ceph/pull/6564> which fixes this, and other issues.

#12 - 11/13/2015 01:08 PM - Nathan Cutler

- *Status changed from In Progress to Need Review*

#13 - 11/17/2015 03:11 PM - Sage Weil

- *Assignee changed from Piotr Dalek to Joao Eduardo Luis*

#14 - 01/07/2016 11:52 AM - Joao Eduardo Luis

- *Category set to Monitor*

- Status changed from Need Review to Resolved