

## Messengers - Bug #40209

### msg: connection thrash failure causes mon to be marked down

06/08/2019 12:07 AM - Patrick Donnelly

<b>Status:</b> New	<b>% Done:</b> 0%
<b>Priority:</b> High	
<b>Assignee:</b>	
<b>Category:</b>	
<b>Target version:</b> v15.0.0	
<b>Source:</b> Q/A	<b>Reviewed:</b>
<b>Tags:</b>	<b>Affected Versions:</b>
<b>Backport:</b> nautilus	<b>ceph-qa-suite:</b>
<b>Regression:</b> No	<b>Pull request ID:</b>
<b>Severity:</b> 3 - minor	<b>Crash signature:</b>
<b>Description</b>	
I would expect this to be recoverable without a health warning. Is that correct?	
Injected failure causing mon.1 (mon.b) to not respond to mon.0 (mon.a):	
<pre>2019-06-07T04:23:07.590+0000 7f5999965700 0 -- [v2:172.21.15.143:3300/0,v1:172.21.15.143:6789/0] &gt;&gt; [v2:172.21.15.44:3300/0,v1:172.21.15.44:6789/0] conn(0x55ec35801c00 msgr2=0x55ec357fcb00 secure :-1 s=STATE_CONNECTION_ESTABLISHED l=0).read_until injecting socket failure</pre>	
<a href="/ceph/teuthology-archive/pdonnell-2019-06-07_01:05:33-fs-wip-pdonnell-testing-20190606.223209-distro-basic-smithi/4009456/remote/smithi143/log/ceph-mon.b.log.gz">/ceph/teuthology-archive/pdonnell-2019-06-07_01:05:33-fs-wip-pdonnell-testing-20190606.223209-distro-basic-smithi/4009456/remote/smithi143/log/ceph-mon.b.log.gz</a>	
<b>Causes this warning:</b>	
<pre>2019-06-07T06:03:48.098 INFO:teuthology.orchestra.run.smithi044.stdout:2019-06-07T04:23:22.623678+ 0000 mon.a (mon.0) 1918 : cluster [WRN] Health check failed: 1/3 mons down, quorum a,c (MON_DOWN)</pre>	
<b>From:</b>	
<a href="/ceph/teuthology-archive/pdonnell-2019-06-07_01:05:33-fs-wip-pdonnell-testing-20190606.223209-distro-basic-smithi/4009456/teuthology.log">/ceph/teuthology-archive/pdonnell-2019-06-07_01:05:33-fs-wip-pdonnell-testing-20190606.223209-distro-basic-smithi/4009456/teuthology.log</a>	
Shouldn't the connection be reestablished transparently?	

## History

### #1 - 06/12/2019 02:16 PM - Patrick Donnelly

More failures:

```
Failure: "2019-06-11T16:47:23.746936+0000 mon.b (mon.0) 955 : cluster [WRN] Health check failed: 1/3 mons down
, quorum b,a (MON_DOWN)" in cluster log
2 jobs: ['4020203', '4020067']
suites intersection: ['conf/{client.yaml', 'mds.yaml', 'mon.yaml', 'mount/fuse.yaml', 'osd.yaml}', 'overrides/
{frag_enable.yaml', 'supported-random-distros$/{ubuntu_latest.yaml}', 'whitelist_health.yaml', 'whitelist_wron
gly_marked_down.yaml}']
suites union: ['ceph-thrash/default.yaml', 'clusters/1-mds-1-client-coloc.yaml', 'clusters/1-mds-4-client-colo
c.yaml', 'conf/{client.yaml', 'fs/basic_functional/{begin.yaml', 'fs/thrash/{begin.yaml', 'mds.yaml', 'mon.yam
l', 'mount/fuse.yaml', 'msgr-failures/osd-mds-delay.yaml', 'no_client_pidfile.yaml', 'objectstore-ec/bluestore
-comp.yaml', 'objectstore/bluestore-bitmap.yaml', 'osd.yaml}', 'overrides/{frag_enable.yaml', 'supported-rando
m-distros$/{ubuntu_latest.yaml}', 'tasks/cfuse_workunit_snaptests.yaml}', 'tasks/client-recovery.yaml}', 'whit
elist_health.yaml', 'whitelist_wrongly_marked_down.yaml}']
```

From: /ceph/teuthology-archive/pdonnell-2019-06-11\_01:05:56-fs-wip-pdonnell-testing-20190610.220401-distro-basic-smithi