

RADOS - Bug #36767

OSD: unrecoverable heartbeat connections

11/12/2018 03:09 PM - Yury Z

Status:	Fix Under Review	% Done:	0%
Priority:	Normal	Spent time:	0.00 hour
Assignee:			
Category:			
Target version:			
Source:		Affected Versions:	v12.2.8
Tags:		ceph-qa-suite:	rados
Backport:		Component(RADOS):	OSD
Regression:	No	Pull request ID:	
Severity:	3 - minor	Crash signature:	
Reviewed:			

Description

There are several unrecoverable heartbeat connections according to logs.
They usually appears after problems/reproducing problems in cluster network.

```
Nov 06 15:30:54 P20B-SR4-R1-CEPH-DS03 ceph-osd[3312534]: 2018-11-06 15:30:54.880354 7f9aae204700 -1 osd.101 298277
heartbeat_check: no reply from 10.8.0.1:6804 osd.129 ever on either front[unhealthy] or back[healthy], first ping sent 2018-11-06
15:25:10.089267 (cutoff 2018-11-06 15:30:44.880331)
```

```
Nov 06 15:30:54 P20B-SR4-R1-CEPH-DS03 ceph-osd[3312534]: 2018-11-06 15:30:54.880354 7f9aae204700 -1 osd.101 298277
heartbeat_check: no reply from 10.8.0.1:6804 osd.129 ever on either front[unhealthy] or back[healthy], first ping sent 2018-11-06
15:25:10.089267 (cutoff 2018-11-06 15:30:44.880331)
```

```
Nov 06 15:30:55 P20B-SR4-R1-CEPH-DS03 ceph-osd[3312534]: 2018-11-06 15:30:55.880650 7f9aae204700 -1 osd.101 298277
heartbeat_check: no reply from 10.8.0.1:6804 osd.129 ever on either front[unhealthy] or back[healthy], first ping sent 2018-11-06
15:25:10.089267 (cutoff 2018-11-06 15:30:45.880639)
```

```
Nov 06 15:30:55 P20B-SR4-R1-CEPH-DS03 ceph-osd[3312534]: 2018-11-06 15:30:55.880650 7f9aae204700 -1 osd.101 298277
heartbeat_check: no reply from 10.8.0.1:6804 osd.129 ever on either front[unhealthy] or back[healthy], first ping sent 2018-11-06
15:25:10.089267 (cutoff 2018-11-06 15:30:45.880639)
```

Messages about unhealthy heartbeat connection doesn't dissappear until osd reboot.

```
Nov 06 15:45:19 P20B-SR4-R1-CEPH-DS03 ceph-osd[3312534]: 2018-11-06 15:45:19.463534 7f9aae204700 -1 osd.101 298278
heartbeat_check: no reply from 10.8.0.1:6804 osd.129 ever on either front[unhealthy] or back[healthy], first ping sent 2018-11-06
15:25:10.089267 (cutoff 2018-11-06 15:45:09.463523)
```

```
Nov 06 15:45:19 P20B-SR4-R1-CEPH-DS03 ceph-osd[3312534]: 2018-11-06 15:45:19.463534 7f9aae204700 -1 osd.101 298278
heartbeat_check: no reply from 10.8.0.1:6804 osd.129 ever on either front[unhealthy] or back[healthy], first ping sent 2018-11-06
15:25:10.089267 (cutoff 2018-11-06 15:45:09.463523)
```

According to tcpdump, client tries to reconnect every 15 seconds and restore heartbeat connection, but server disconnects it.

```
15:40:07.854454 IP 10.8.0.8.39996 > 10.8.0.1.6804: Flags [S], seq 3822599484, win 26880, options [mss 8960,sackOK,TS val
85988139 ecr 0,nop,wscale 7], length 0
```

```
15:40:07.854492 IP 10.8.0.1.6804 > 10.8.0.8.39996: Flags [S.], seq 1273155899, ack 3822599485, win 26844, options [mss
8960,sackOK,TS val 205723552 ecr 85988139,nop,wscale 7], length 0
```

```
15:40:07.854663 IP 10.8.0.8.39996 > 10.8.0.1.6804: Flags [.], ack 1, win 210, options [nop,nop,TS val 85988140 ecr 205723552],
length 0
```

```
15:40:07.854729 IP 10.8.0.8.39996 > 10.8.0.1.6804: Flags [P.], seq 1:10, ack 1, win 210, options [nop,nop,TS val 85988140 ecr
205723552], length 9
```

```
15:40:07.854738 IP 10.8.0.1.6804 > 10.8.0.8.39996: Flags [.], ack 10, win 210, options [nop,nop,TS val 205723552 ecr 85988140],
length 0
```

```
15:40:07.855275 IP 10.8.0.1.6804 > 10.8.0.8.39996: Flags [P.], seq 1:282, ack 10, win 210, options [nop,nop,TS val 205723552 ecr
85988140], length 281
```

```
15:40:07.855412 IP 10.8.0.8.39996 > 10.8.0.1.6804: Flags [.], ack 282, win 219, options [nop,nop,TS val 85988140 ecr 205723552],
length 0
```

```
15:40:07.855454 IP 10.8.0.8.39996 > 10.8.0.1.6804: Flags [P.], seq 10:146, ack 282, win 219, options [nop,nop,TS val 85988140 ecr
```

205723552], length 136
15:40:07.855461 IP 10.8.0.8.39996 > 10.8.0.1.6804: Flags [P.], seq 146:179, ack 282, win 219, options [nop,nop,TS val 85988140 ecr 205723552], length 33
15:40:07.855489 IP 10.8.0.1.6804 > 10.8.0.8.39996: Flags [.] , ack 179, win 219, options [nop,nop,TS val 205723553 ecr 85988140], length 0
15:40:07.855544 IP 10.8.0.1.6804 > 10.8.0.8.39996: Flags [F.], seq 282, ack 179, win 219, options [nop,nop,TS val 205723553 ecr 85988140], length 0
15:40:07.855722 IP 10.8.0.8.39996 > 10.8.0.1.6804: Flags [F.], seq 179, ack 283, win 219, options [nop,nop,TS val 85988141 ecr 205723553], length 0
15:40:07.855730 IP 10.8.0.1.6804 > 10.8.0.8.39996: Flags [.] , ack 180, win 219, options [nop,nop,TS val 205723553 ecr 85988141], length 0

Server closes connection because there is race with existing connection.
Existing connection already stopped, but not completely removed.

Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855243 7fc036a77700 20 -- 10.8.0.1:6804/12800635 >> - conn(0x55b30b151000 :-1 s=STATE_ACCEPTING pgs=0 cs=0 l=0).process prev state is STATE_ACCEPTING
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855261 7fc036a77700 1 -- 10.8.0.1:6804/12800635 >> - conn(0x55b30b151000 :6804 s=STATE_ACCEPTING pgs=0 cs=0 l=0)._process_connection sd=558 -
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855280 7fc036a77700 10 -- 10.8.0.1:6804/12800635 >> - conn(0x55b30b151000 :6804 s=STATE_ACCEPTING pgs=0 cs=0 l=0)._try_send sent bytes 281 remaining bytes 0
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855282 7fc036a77700 10 -- 10.8.0.1:6804/12800635 >> - conn(0x55b30b151000 :6804 s=STATE_ACCEPTING_WAIT_BANNER_ADDR pgs=0 cs=0 l=0)._process_connection write banner and addr done: -
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855284 7fc036a77700 20 -- 10.8.0.1:6804/12800635 >> - conn(0x55b30b151000 :6804 s=STATE_ACCEPTING_WAIT_BANNER_ADDR pgs=0 cs=0 l=0).process prev state is STATE_ACCEPTING
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855464 7fc036a77700 20 -- 10.8.0.1:6804/12800635 >> - conn(0x55b30b151000 :6804 s=STATE_ACCEPTING_WAIT_BANNER_ADDR pgs=0 cs=0 l=0).process prev state is STATE_ACCEPTING_WAIT_BANNER_ADDR
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855493 7fc036a77700 10 -- 10.8.0.1:6804/12800635 >> - conn(0x55b30b151000 :6804 s=STATE_ACCEPTING_WAIT_BANNER_ADDR pgs=0 cs=0 l=0)._process_connection accept peer addr is 10.8.0.8:0/3312534
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855496 7fc036a77700 20 -- 10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804 s=STATE_ACCEPTING_WAIT_CONNECT_MSG pgs=0 cs=0 l=0).process prev state is STATE_ACCEPTING_WAIT_BANNER_ADDR
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855499 7fc036a77700 20 -- 10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804 s=STATE_ACCEPTING_WAIT_CONNECT_MSG_AUTH pgs=0 cs=0 l=0).process prev state is STATE_ACCEPTING_WAIT_CONNECT_MSG
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855500 7fc036a77700 20 -- 10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804 s=STATE_ACCEPTING_WAIT_CONNECT_MSG_AUTH pgs=0 cs=0 l=0)._process_connection accept got peer connect_seq 0 global_seq 1709
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855503 7fc036a77700 10 -- 10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804 s=STATE_ACCEPTING_WAIT_CONNECT_MSG_AUTH pgs=0 cs=0 l=1)._process_connection accept of host_type 4, policy.lossy=1 policy.server=1 policy.standby=0 policy.resetcheck=0
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855505 7fc036a77700 10 -- 10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804 s=STATE_ACCEPTING_WAIT_CONNECT_MSG_AUTH pgs=0 cs=0 l=1).handle_connect_msg accept my proto 10, their proto 10
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855508 7fc036a77700 10 -- 10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804 s=STATE_ACCEPTING_WAIT_CONNECT_MSG_AUTH pgs=0 cs=0 l=1).handle_connect_msg accept setting up session_security.
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855511 7fc036a77700 1 -- 10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804 s=STATE_ACCEPTING_WAIT_CONNECT_MSG_AUTH pgs=0 cs=0 l=1).handle_connect_msg existing already closed.
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855513 7fc036a77700 10 -- 10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804 s=STATE_ACCEPTING_WAIT_CONNECT_MSG_AUTH pgs=1709 cs=1 l=1).handle_connect_msg accept success, connect_seq = 1 in_seq=0, sending READY
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855515 7fc036a77700 10 -- 10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804 s=STATE_ACCEPTING_WAIT_CONNECT_MSG_AUTH pgs=1709 cs=1 l=1).handle_connect_msg accept features 4611087853745930235

```
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855518 7fc036a77700 1 --
10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804
s=STATE_ACCEPTING_WAIT_CONNECT_MSG_AUTH pgs=1709 cs=1 l=1).handle_connect_msg existing race replacing process
for addr=10.8.0.8:0/3312534 just fail later one(this)
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855521 7fc036a77700 10 --
10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804
s=STATE_ACCEPTING_WAIT_CONNECT_MSG_AUTH pgs=1709 cs=1 l=1).handle_connect_msg accept fault after register
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855522 7fc036a77700 10 --
10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804
s=STATE_ACCEPTING_WAIT_CONNECT_MSG_AUTH pgs=1709 cs=1 l=1).handle_connect_msg failed to accept.
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855524 7fc036a77700 1 --
10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804
s=STATE_ACCEPTING_WAIT_CONNECT_MSG_AUTH pgs=1709 cs=1 l=1).fault on lossy channel, failing
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855526 7fc036a77700 2 --
10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804
s=STATE_ACCEPTING_WAIT_CONNECT_MSG_AUTH pgs=1709 cs=1 l=1)._stop
Nov 06 15:40:07 P20B-SR4-R2-CEPH-DS01 ceph-osd[2800635]: 2018-11-06 15:40:07.855528 7fc036a77700 10 --
10.8.0.1:6804/12800635 >> 10.8.0.8:0/3312534 conn(0x55b30b151000 :6804
s=STATE_ACCEPTING_WAIT_CONNECT_MSG_AUTH pgs=1709 cs=1 l=1).discard_out_queue started
```

History

#1 - 11/12/2018 03:41 PM - Yury Z

Pull request:

<https://github.com/ceph/ceph/pull/25061>

#2 - 11/16/2018 04:53 AM - Kefu Chai

- Status changed from New to Fix Under Review

#3 - 11/22/2018 07:32 AM - Yan Jun

see also: <https://tracker.ceph.com/issues/36175>