

## Ceph - Bug #2116

### Repeated messages of "heartbeat\_check: no heartbeat from"

02/28/2012 07:35 AM - Wido den Hollander

<b>Status:</b>	Resolved	<b>Start date:</b>	02/28/2012
<b>Priority:</b>	High	<b>Due date:</b>	
<b>Assignee:</b>		<b>% Done:</b>	0%
<b>Category:</b>	OSD	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	v0.44	<b>Spent time:</b>	18.00 hours
<b>Source:</b>	Community (dev)	<b>Reviewed:</b>	
<b>Tags:</b>		<b>Affected Versions:</b>	
<b>Backport:</b>		<b>ceph-qa-suite:</b>	
<b>Regression:</b>	No	<b>Pull request ID:</b>	
<b>Severity:</b>	3 - minor		

#### Description

As discussed on the ml I gathered some logs.

Today I upgraded my whole cluster to 0.42.2 from 0.41.

Due to the on-disk change I formatted my cluster and started it again.

Immediately after the start of the OSD's the "no heartbeat from" started:

```
2012-02-28 16:30:31.951132    pg v708: 7920 pgs: 7920 active+clean; 8730 bytes data, 164 MB used,
74439 GB / 74520 GB avail
2012-02-28 16:30:31.980965    mds e4: 1/1/1 up {0=alpha=up:active}
2012-02-28 16:30:31.981010    osd e65: 40 osds: 40 up, 40 in
2012-02-28 16:30:31.981192    log 2012-02-28 16:30:30.344847 mon.0 [2a00:f10:11a:408::1]:6789/0 107
28 : [INF] osd.5 [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6803/14395 failed (by osd.15 [2a00:f10:11b
:cef0:225:90ff:fe33:49b0]:6809/12129)
2012-02-28 16:30:31.981314    mon e1: 3 mons at {pri=[2a00:f10:11b:cef0:230:48ff:fed3:b086]:6789/0,
sec=[2a00:f10:11a:408::1]:6789/0,third=[2a00:f10:11a:409::1]:6789/0}
2012-02-28 16:30:33.680286    log 2012-02-28 16:30:33.533593 mon.0 [2a00:f10:11a:408::1]:6789/0 107
29 : [INF] osd.0 [2a00:f10:11b:cef0:225:90ff:fe33:49fe]:6800/12003 failed (by osd.13 [2a00:f10:11b
:cef0:225:90ff:fe33:49b0]:6803/11917)
2012-02-28 16:30:34.705162    log 2012-02-28 16:30:33.709207 mon.0 [2a00:f10:11a:408::1]:6789/0 107
30 : [INF] osd.6 [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6806/14486 failed (by osd.9 [2a00:f10:11b:
cef0:225:90ff:fe33:49f2]:6809/19176)
2012-02-28 16:30:34.705162    log 2012-02-28 16:30:33.968388 mon.0 [2a00:f10:11a:408::1]:6789/0 107
31 : [INF] osd.30 [2a00:f10:11b:cef0:225:90ff:fe33:499c]:6806/29719 failed (by osd.37 [2a00:f10:11
b:cef0:225:90ff:fe33:49a4]:6803/26068)
2012-02-28 16:30:35.737775    log 2012-02-28 16:30:35.346912 mon.0 [2a00:f10:11a:408::1]:6789/0 107
32 : [INF] osd.5 [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6803/14395 failed (by osd.15 [2a00:f10:11b
:cef0:225:90ff:fe33:49b0]:6809/12129)
```

As you can see, the cluster is completely fresh. No data at all and no I/O load.

One of the things that came to mind was that it might be a clock issue, not being synchronized. I verified all the clocks and those are synchronized.

Information on all the OSD's:

- ceph version 0.42.2 ([732f3ec94e39d458230b7728b2a936d431e19322](https://github.com/ceph/ceph/blob/732f3ec94e39d458230b7728b2a936d431e19322))
- Kernel: 3.2.0
- Memory: 4GB
- CPU: Atom D525
- Disks: 2TB Seagate/WD

## History

### #1 - 02/28/2012 09:22 AM - Sage Weil

- Target version set to v0.44

### #2 - 02/28/2012 11:20 AM - Sage Weil

- Status changed from New to In Progress

looks like a msgr issue?

```
2012-02-28 12:47:45.167980 7fd09f94e700 -- [2a00:f10:11b:cef0:225:90ff:fe33:49b0]:0/12130 --> [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6805/14395 -- osd_ping(start_heartbeat e0 as_of 20) v1 -- ?+0 0x237cc40 con 0x2181500
2012-02-28 12:47:45.839282 7fd09f94e700 -- [2a00:f10:11b:cef0:225:90ff:fe33:49b0]:0/12130 --> [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6805/14395 -- osd_ping(start_heartbeat e0 as_of 25) v1 -- ?+0 0x237cc40 con 0x21cea00
2012-02-28 12:47:45.855370 7fd097b1d700 -- [2a00:f10:11b:cef0:225:90ff:fe33:49b0]:0/12130 >> [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6805/14395 pipe(0x2228500 sd=41 pgs=24 cs=1 l=0).fault initiating reconnect
2012-02-28 12:47:45.871679 7fd098729700 -- [2a00:f10:11b:cef0:225:90ff:fe33:49b0]:0/12130 >> [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6805/14395 pipe(0x2228500 sd=41 pgs=24 cs=2 l=0).connect got RESETSESSION
2012-02-28 13:02:45.941694 7fd095e00700 -- [2a00:f10:11b:cef0:225:90ff:fe33:49b0]:0/12130 >> [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6805/14395 pipe(0x2228500 sd=41 pgs=26 cs=1 l=0).fault with nothing to send, going to st andby
```

### #3 - 02/28/2012 11:26 AM - Sage Weil

the other side of this conversation is

```
2012-02-28 12:47:45.200092 7f7097ca1700 -- [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6805/14395 <== osd.15 [2a00:f10:11b:cef0:225:90ff:fe33:49b0]:0/12130 1 ==== osd_ping(start_heartbeat e0 as_of 20) v1 ==== 39+0+0 (50269974 9 0 0) 0x2ce78c0 con 0x26a2780
2012-02-28 12:47:45.228604 7f7090064700 -- [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6805/14395 >> [2a00:f10:11b:cef0:225:90ff:fe33:49b0]:0/12130 pipe(0x274ea00 sd=44 pgs=11 cs=1 l=0).fault with nothing to send, going to st andby
2012-02-28 12:47:45.630629 7f7093b98700 -- [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6805/14395 --> [2a00:f10:11b:cef0:225:90ff:fe33:49b0]:0/12130 -- osd_ping(heartbeat e0 as_of 20) v1 -- ?+0 0x2d60e00 con 0x26a2780
2012-02-28 12:47:45.847321 7f708eb4f700 -- [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6805/14395 >> [2a00:f10:11b:cef0:225:90ff:fe33:49b0]:0/12130 pipe(0x282b000 sd=60 pgs=0 cs=0 l=0).accept connect_seq 0 vs existing 2 state 1
2012-02-28 12:47:45.847411 7f708eb4f700 -- [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6805/14395 >> [2a00:f10:11b:cef0:225:90ff:fe33:49b0]:0/12130 pipe(0x282b000 sd=60 pgs=0 cs=0 l=0).accept peer reset, then tried to connect to us, replacing
2012-02-28 12:47:45.848900 7f7097ca1700 -- [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6805/14395 <== osd.15 [2a00:f10:11b:cef0:225:90ff:fe33:49b0]:0/12130 1 ==== osd_ping(start_heartbeat e0 as_of 25) v1 ==== 39+0+0 (24242950 25 0 0) 0x2d60e00 con 0x2e8ec80
2012-02-28 12:47:45.870002 7f708eb4f700 -- [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6805/14395 >> [2a00:f10:11b:cef0:225:90ff:fe33:49b0]:0/12130 pipe(0x282b280 sd=56 pgs=0 cs=0 l=0).accept we reset (peer sent cseq 2), sending RESETSESSION
2012-02-28 13:02:45.940175 7f708eb4f700 -- [2a00:f10:11b:cef0:225:90ff:fe32:cf64]:6805/14395 >> [2a00:f10:11b:cef0:225:90ff:fe33:49b0]:0/12130 pipe(0x282b280 sd=56 pgs=24 cs=1 l=0).fault with nothing to send, going to st andby
```

#### #4 - 02/28/2012 04:23 PM - Greg Farnum

This may be a messenger issue, but it's not losing that initial message — notice how osd5 tries to send a ping back to osd5? For some reason it doesn't apparently trigger a reconnect though, which is odd...  
Also the bit where osd5 is apparently getting out-of-order connects...

#### #5 - 02/29/2012 02:31 PM - Sage Weil

Wido, are you able to reproduce this reliably? I have an idea what the problem is, but have never reproduced this. Hopefully you can verify it goes away?

#### #6 - 02/29/2012 02:57 PM - Sage Weil

- Status changed from In Progress to Feedback

i'm hoping wip-2116 fixes it...

#### #7 - 03/01/2012 02:48 AM - Wido den Hollander

- File *osd.3.heartbeat.log.gz* added

- File *osd.8.heartbeat.log.gz* added

I can almost always reproduce it.

I just upgraded my cluster to:

ceph version 0.42.2-206-gd77c579 (commit:d77c5796fc898c8fdf6b38fe887c5c62e366023e)

It all worked fine for about 20 minutes and then I started seeing this again:

```
2012-03-01 11:44:10.083448 log 2012-03-01 11:44:09.957804 mon.0 [2a00:f10:11a:408::1]:6789/0 132126 : [INF]
osd.8 [2a00:f10:11b:cef0:225:90ff:fe33:49f2]:6818/4327 failed (by osd.3 [2a00:f10:11b:cef0:225:90ff:fe33:49fe]
:6804/20053)
2012-03-01 11:44:10.083448 log 2012-03-01 11:44:09.958113 mon.0 [2a00:f10:11a:408::1]:6789/0 132127 : [INF]
osd.39 [2a00:f10:11b:cef0:225:90ff:fe33:49a4]:6820/11945 failed (by osd.3 [2a00:f10:11b:cef0:225:90ff:fe33:49fe]
e]:6804/20053)
```

I've attached the logs of osd.0 and osd.8 again, hope that helps.

```
2012-03-01 11:46:09.729570 7f14027af700 osd.3 360 heartbeat_check: no heartbeat from osd.8 since 2012-03-01 10
:51:46.452920 (cutoff 2012-03-01 11:45:49.729539)
2012-03-01 11:46:10.014949 7f140f8ca700 osd.3 360 heartbeat_check: no heartbeat from osd.8 since 2012-03-01 10
:51:46.452920 (cutoff 2012-03-01 11:45:50.014920)
2012-03-01 11:46:10.015179 7f140f8ca700 -- [2a00:f10:11b:cef0:225:90ff:fe33:49fe]:6804/20053 --> [2a00:f10:11b
:cef0:230:48ff:fed3:b086]:6789/0 -- osd_failure(osd.8 [2a00:f10:11b:cef0:225:90ff:fe33:49f2]:6818/4327 e360 v3
60) v2 -- ?+0 0x269d000 con 0x1894780
2012-03-01 11:46:10.432683 7f14027af700 osd.3 360 heartbeat_check: no heartbeat from osd.8 since 2012-03-01 10
:51:46.452920 (cutoff 2012-03-01 11:45:50.432651)
```

I'll leave the cluster running like this, that will probably keep it in this state. If any further logging is needed I can turn that on.

**#8 - 03/02/2012 03:08 PM - Sage Weil**

ok, i have a theory what's going on. can you try the new wip-2116, and run with debug ms = 20?

thanks!

**#9 - 03/05/2012 09:48 AM - Sage Weil**

- Priority changed from Normal to High

**#10 - 03/05/2012 01:23 PM - Wido den Hollander**

I tried to compile the wip-2116 and master branches today, that failed. So I haven't been able to test the new code yet.

I'll try again tomorrow and see if I get it working.

**#11 - 03/06/2012 08:25 AM - Wido den Hollander**

I ran multiple I/O loads on the cluster last night and performed multiple tests, but this problem didn't come back yet.

I wouldn't close this issue yet, let's wait a couple of days to see if it might come back.

**#12 - 03/09/2012 10:58 AM - Wido den Hollander**

I'm not willing to say it's completely gone, but it is starting to seem like it.

In the past few days the cluster has been running fine with a moderate I/O load, no issues at all.

For now we could close this issue, I'll re-open it when I see it again.

**#13 - 03/09/2012 11:07 AM - Sage Weil**

Hrm... well, if that's the case then we should merge the wip-2116 as is, then... putting that on my list

**#14 - 03/22/2012 09:14 AM - Sage Weil**

see new wip-osd-hb branch

**#15 - 03/30/2012 03:59 PM - Sage Weil**

- Status changed from Feedback to Resolved

fixed by [374bef9c97266600b4c6b83100485d7250363213](https://github.com/sage/weil/commit/374bef9c97266600b4c6b83100485d7250363213)

**Files**

---

ceph.conf	6.67 KB	02/28/2012	Wido den Hollander
osd.5.log_heartbeat.gz	28.8 MB	02/28/2012	Wido den Hollander
osd.15.log_heartbeat.gz	25.5 MB	02/28/2012	Wido den Hollander
osd.3.heartbeat.log.gz	32.2 MB	03/01/2012	Wido den Hollander
osd.8.heartbeat.log.gz	26.2 MB	03/01/2012	Wido den Hollander