

## Ceph - Bug #14175

### clock skew report is incorrect by "ceph health detail" command

12/24/2015 05:59 AM - wei qiaomiao

<b>Status:</b>	Resolved	<b>Start date:</b>	12/24/2015
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>	Joao Eduardo Luis	<b>% Done:</b>	0%
<b>Category:</b>	Monitor	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>		<b>Spent time:</b>	0.00 hour
<b>Source:</b>	other	<b>Reviewed:</b>	
<b>Tags:</b>		<b>Affected Versions:</b>	
<b>Backport:</b>	hammer	<b>ceph-qa-suite:</b>	
<b>Regression:</b>	No	<b>Pull request ID:</b>	
<b>Severity:</b>	3 - minor		
<b>Description</b>			
I use "ceph health detail" command to check my cluster health and found below warning:  mon.c7 addr 10.118.202.97:6789/0 clock skew 239.478s > max 0.05s (latency 0.0355416s)  so i modify mon.c7 system time to make it the same as the leader monitor, but the warning is still exist:  mon.c7 addr 10.118.202.97:6789/0 clock skew 191.582s > max 0.05s (latency 0.0286543s)3s			
<b>Related issues:</b>			
Copied to Ceph - Backport #15024: hammer: clock skew report is incorrect by "...			<b>Resolved</b>

#### Associated revisions

##### Revision 17d8ff42 - 03/03/2016 11:06 AM - Joao Eduardo Luis

mon: Monitor: get rid of weighted clock skew reports

By weighting the reports we were making it really hard to get rid of a clock skew warning once the cause had been fixed.

Instead, as soon as we get a clean bill of health, let's run a new round and soon as possible and ascertain whether that was a transient fix or for realies. That should be better than the alternative of waiting for an hour or something (for a large enough skew) for the warning to go away - and with it, the admin's sanity ("WHAT AM I DOING WRONG???").

Fixes: #14175

Signed-off-by: Joao Eduardo Luis <[joao@suse.de](mailto:joao@suse.de)>

##### Revision 01672b4a - 03/14/2016 12:50 AM - Joao Eduardo Luis

mon: Monitor: get rid of weighted clock skew reports

By weighting the reports we were making it really hard to get rid of a

clock skew warning once the cause had been fixed.

Instead, as soon as we get a clean bill of health, let's run a new round and soon as possible and ascertain whether that was a transient fix or for realies. That should be better than the alternative of waiting for an hour or something (for a large enough skew) for the warning to go away - and with it, the admin's sanity ("WHAT AM I DOING WRONG???").

Fixes: #14175

Signed-off-by: Joao Eduardo Luis <[joao@suse.de](mailto:joao@suse.de)>

(cherry pick from commit 17d8ff429c7dca8fc1ada6e7cc8a7c4924a22e28)

## History

---

### #1 - 12/24/2015 12:51 PM - Nathan Cutler

- Assignee set to Joao Eduardo Luis

### #2 - 12/24/2015 01:01 PM - Nathan Cutler

Hopefully Joao will chime in with a deeper explanation, but until then I can say that I have run into a similar issue (you don't mention which Ceph version you are using - I was using 0.94.3).

Here is what I remember of Joao's explanation:

As of Hammer there is new clock-skew handling logic in the monitor code that is designed to make the cluster more tolerant of time discrepancies (situations where the clock on one monitor node is slightly ahead of, or behind, the other nodes). This, however, comes with an unintended side effect: it now takes longer for clusters to recover from large time differences.

Whether or not this is a bug is still an open question.

It would be interesting to know how long it takes the cluster to recover from clock skew reported here. In my case, the time discrepancy was arising at boot time and the cluster took 15-60 minutes to recover (i.e. for the "clock skew" warning to disappear).

### #3 - 12/25/2015 06:25 AM - wei qiaomiao

I was using 0.94.5 version

How long it take the cluster recover fron clock skew reported depend on how large time of the clock drift. In my environment the cluster took 3-4 hours to recover when clock drift is 2-3 minutes. It's too long time for user

May be we can improve the clock-skew handling mechanism for the scene of cluster's clock drift is large, for example, when the absolute of current clock skew value minus the last value is larger than 5s(or other value we can discuss), we drop the last value and only took the current value to report

**#4 - 01/07/2016 11:36 AM - Joao Eduardo Luis**

- *Category set to Monitor*
- *Status changed from New to Need Review*

<https://github.com/ceph/ceph/pull/7141>

**#5 - 03/08/2016 04:32 PM - Kefu Chai**

- *Status changed from Need Review to Pending Backport*

Joao, shall we backport this change to hammer?

**#6 - 03/08/2016 07:41 PM - Nathan Cutler**

- *Backport set to hammer*

**#7 - 03/09/2016 03:27 AM - Nathan Cutler**

- *Copied to Backport #15024: hammer: clock skew report is incorrect by "ceph health detail" command added*

**#8 - 05/16/2016 10:25 AM - Loic Dachary**

- *Status changed from Pending Backport to Resolved*