

RADOS - Bug #55383

monitor cluster logs(ceph.log) appear empty until rotated

04/19/2022 04:43 PM - Tim Wilkinson

Status:	Resolved	% Done:	0%
Priority:	Immediate	Spent time:	0.00 hour
Assignee:			
Category:	Monitor		
Target version:			
Source:		Affected Versions:	v17.0.0
Tags:		ceph-qa-suite:	rados
Backport:	quincy	Component(RADOS):	
Regression:	No	Pull request ID:	46124
Severity:	3 - minor	Crash signature (v1):	
Reviewed:		Crash signature (v2):	
Description			
While executing tests on 17.1.0-203-g2c8d01fc, I see the ceph.log files on all MONs are zero length unless rotated.			
<pre># ansible mons -m shell -a "ls -l /var/log/ceph/78fbaf12-bea9-11ec-8fd2-000af7995756/ceph.log*" f28-h29-000-r630 CHANGED rc=0 >> -rw----- 1 ceph ceph 0 Apr 19 03:18 /var/log/ceph/78fbaf12-bea9-11ec-8fd2-000af7995756/*ceph.log* -rw----- 1 ceph ceph 1359181 Apr 19 03:18 /var/log/ceph/78fbaf12-bea9-11ec-8fd2-000af7995756/ceph.log-20220419.gz f28-h30-000-r630 CHANGED rc=0 >> -rw----- 1 ceph ceph 0 Apr 19 03:08 /var/log/ceph/78fbaf12-bea9-11ec-8fd2-000af7995756/*ceph.log* -rw----- 1 ceph ceph 1350438 Apr 19 03:07 /var/log/ceph/78fbaf12-bea9-11ec-8fd2-000af7995756/ceph.log-20220419.gz f28-h28-000-r630 CHANGED rc=0 >> -rw----- 1 ceph ceph 0 Apr 19 03:43 /var/log/ceph/78fbaf12-bea9-11ec-8fd2-000af7995756/*ceph.log* -rw----- 1 ceph ceph 1380492 Apr 19 03:42 /var/log/ceph/78fbaf12-bea9-11ec-8fd2-000af7995756/ceph.log-20220419.gz</pre>			
Related issues:			
Copied to RADOS - Backport #55742: quincy: monitor cluster logs(ceph.log) app...			Resolved

History

#1 - 04/19/2022 04:49 PM - Vikhyat Umrao

- Subject changed from monitor logs appear empty until rotated to monitor cluster logs(ceph.log) appear empty until rotated

#2 - 04/19/2022 04:51 PM - Vikhyat Umrao

Tim Wilkinson wrote:

While executing tests on 17.1.0-203-g2c8d01fc, I see the ceph.log files on all MONs are zero length unless rotated.

[...]

To me it looks like after log rotation, logging in the cluster logs stops? I can see the same in the Gibba cluster.

```
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ls -ltr ceph.log*
-rw----- . 1 ceph ceph 1350774 Apr  5 03:33 ceph.log-20220405.gz
-rw----- . 1 ceph ceph 297426 Apr  6 03:46 ceph.log-20220406.gz
```

```
-rw-----. 1 ceph ceph 141884 Apr 8 03:07 ceph.log-20220408.gz
-rw-----. 1 ceph ceph 678337 Apr 9 03:34 ceph.log-20220409.gz
-rw-----. 1 ceph ceph 389794 Apr 12 03:09 ceph.log-20220412.gz
-rw-----. 1 ceph ceph 1241977 Apr 13 03:26 ceph.log-20220413.gz
-rw-----. 1 ceph ceph 788096 Apr 14 03:24 ceph.log-20220414.gz
-rw-----. 1 ceph ceph 0 Apr 14 03:24 ceph.log
```

#3 - 04/19/2022 04:52 PM - Vikhyat Umrao

I think we have not seen this issue before in previous quincy builds? something changed recently.

#4 - 04/19/2022 04:56 PM - Vikhyat Umrao

Gibba cluster quincy version `17.1.0-163-g4e244311`.

#5 - 04/19/2022 05:48 PM - Vikhyat Umrao

Tim Wilkinson wrote:

While executing tests on 17.1.0-203-g2c8d01fc, I see the ceph.log files on all MONs are zero length unless rotated.

[...]

```
root@f28-h28-000-r630:/var/log/ceph/78fbaf12-bea9-11ec-8fd2-000af7995756
# ls -ltr ceph.log*
-rw----- 1 ceph ceph 1380492 Apr 19 03:42 ceph.log-20220419.gz
-rw----- 1 ceph ceph 0 Apr 19 03:43 ceph.log
```

```
root@f28-h28-000-r630:/var/log/ceph/78fbaf12-bea9-11ec-8fd2-000af7995756
# ceph config dump | grep log_to
global                                basic      log_to_file                                true
global                                advanced  mon_cluster_log_to_file                    true
```

Log rotation line in mon log:

```
root@f28-h28-000-r630:/var/log/ceph/78fbaf12-bea9-11ec-8fd2-000af7995756
# zgrep kill ceph-mon.f28-h28-000-r630.log-20220419.gz
2022-04-19T03:43:01.456+0000 7f9648ff2700 -1 received signal: Hangup from Kernel ( Could be generated by pthread_kill(), raise(), abort(), alarm() ) UID: 0
```

#6 - 04/19/2022 11:45 PM - Vikhyat Umrao

- Project changed from RADOS to Orchestrator

- Category set to cephadm

we found out that quincy has <https://github.com/ceph/ceph/pull/40640> log_to_journald feature. When we set

```
ceph config set global mon_cluster_log_to_journald false
```

It fixes the issue and ceph.log, ceph.audit.log logging starts working. So we have a workaround for now but we should still verify why for cluster logs this journald option needs to be disabled

We think this log_to_journald feature should go to quincy Release Notes to avoid confusion for end-users.

#7 - 04/20/2022 08:22 PM - Vikhyat Umrao

Neha noticed today that in the LRC cluster after having this workaround still, this cluster when went through a log rotation window again cluster logs(ceph.log) have stopped updating. See following no ceph.log file.

```
root@reesi004:/var/log/ceph/28f7427e-5558-4ffd-aela-51ec3042759a# ls -lrt *ceph.log*
-rw----- 1 167 167 252931 Feb 28 00:24 ceph.log.7.gz
-rw----- 1 167 167 257361 Feb 28 06:25 ceph.log.6.gz
-rw----- 1 167 167 250446 Feb 28 12:25 ceph.log.5.gz
-rw----- 1 167 167 501828 Feb 28 18:25 ceph.log.4.gz
-rw----- 1 167 167 494610 Apr 18 18:25 ceph.log.3.gz
-rw----- 1 167 167 50904 Apr 20 00:24 ceph.log.2.gz
-rw----- 1 167 167 243892 Apr 20 06:25 ceph.log.1.gz
```

- Then I verified the same in the gibba cluster and it has the similar issue after having the given workaround in place here also ceph.log is not updating after the log rotation window.

```
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ls -lrt ceph.log*
-rw----- 1 ceph ceph 297426 Apr 6 03:46 ceph.log-20220406.gz
-rw----- 1 ceph ceph 141884 Apr 8 03:07 ceph.log-20220408.gz
-rw----- 1 ceph ceph 678337 Apr 9 03:34 ceph.log-20220409.gz
-rw----- 1 ceph ceph 389794 Apr 12 03:09 ceph.log-20220412.gz
-rw----- 1 ceph ceph 1241977 Apr 13 03:26 ceph.log-20220413.gz
-rw----- 1 ceph ceph 788096 Apr 14 03:24 ceph.log-20220414.gz
-rw----- 1 ceph ceph 179467 Apr 20 03:35 ceph.log-20220420.gz
-rw----- 1 ceph ceph 0 Apr 20 03:35 ceph.log
```

```
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ceph config dump | grep journ
global          advanced  mon_cluster_log_to_journald          false
```

- Then just for verification I just did a rm of the workaround option to see if it fixes the issue. And it did fix the issue, it looks like the behavior we see here is very much directly related to this option - mon_cluster_log_to_journald and this journald_log feature came with this PR - <https://github.com/ceph/ceph/pull/40640>.

```
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ls -ltr ceph.log*
-rw----- 1 ceph ceph 297426 Apr 6 03:46 ceph.log-20220406.gz
-rw----- 1 ceph ceph 141884 Apr 8 03:07 ceph.log-20220408.gz
-rw----- 1 ceph ceph 678337 Apr 9 03:34 ceph.log-20220409.gz
-rw----- 1 ceph ceph 389794 Apr 12 03:09 ceph.log-20220412.gz
```

```
-rw-----. 1 ceph ceph 1241977 Apr 13 03:26 ceph.log-20220413.gz
-rw-----. 1 ceph ceph 788096 Apr 14 03:24 ceph.log-20220414.gz
-rw-----. 1 ceph ceph 179467 Apr 20 03:35 ceph.log-20220420.gz
-rw-----. 1 ceph ceph 0 Apr 20 03:35 ceph.log
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ceph config dump | grep journ
global          advanced  mon_cluster_log_to_journald      false
```

```
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ceph config rm global mon_cluster_log_to_journald
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ceph config dump | grep journ
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ls -ltr ceph.log*
-rw-----. 1 ceph ceph 297426 Apr 6 03:46 ceph.log-20220406.gz
-rw-----. 1 ceph ceph 141884 Apr 8 03:07 ceph.log-20220408.gz
-rw-----. 1 ceph ceph 678337 Apr 9 03:34 ceph.log-20220409.gz
-rw-----. 1 ceph ceph 389794 Apr 12 03:09 ceph.log-20220412.gz
-rw-----. 1 ceph ceph 1241977 Apr 13 03:26 ceph.log-20220413.gz
-rw-----. 1 ceph ceph 788096 Apr 14 03:24 ceph.log-20220414.gz
-rw-----. 1 ceph ceph 179467 Apr 20 03:35 ceph.log-20220420.gz
-rw-----. 1 ceph ceph 1464 Apr 20 17:51 ceph.log
```

- Making changes to this workaround option(set and rm of this option) fixes the issue in all the mons, as can be seen in gibba002 also.

```
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ssh gibba002
Activate the web console with: systemctl enable --now cockpit.socket
```

```
Last login: Tue Apr 19 23:04:47 2022 from 172.21.2.101
[root@gibba002 ~]# cd /var/log/ceph/2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8/
[root@gibba002 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ls -ltr ceph.log*
-rw-----. 1 ceph ceph 2590403 Mar 30 03:23 ceph.log-20220330.gz
-rw-----. 1 ceph ceph 241346 Mar 31 03:32 ceph.log-20220331.gz
-rw-----. 1 ceph ceph 1336038 Apr 5 03:37 ceph.log-20220405.gz
-rw-----. 1 ceph ceph 153658 Apr 8 03:32 ceph.log-20220408.gz
-rw-----. 1 ceph ceph 233324 Apr 9 03:41 ceph.log-20220409.gz
-rw-----. 1 ceph ceph 405996 Apr 12 03:36 ceph.log-20220412.gz
-rw-----. 1 ceph ceph 156751 Apr 20 03:24 ceph.log-20220420.gz
-rw-----. 1 ceph ceph 12198 Apr 20 17:53 ceph.log
[root@gibba002 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]#
```

This workaround is useful but may need a cron job, as it needs daily change after log rotation. The issue needs more investigation, why it is causing an issue, and why this option change is helping it.

#8 - 04/21/2022 07:28 PM - Vikhyat Umrao

- Just for completeness as expected this issue again happed today in gibba cluster during the log rotation window and changing the option `mon_cluster_log_to_journald` fixed the issue.

```
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ls -ltr ceph.log*
-rw-----. 1 ceph ceph 141884 Apr  8 03:07 ceph.log-20220408.gz
-rw-----. 1 ceph ceph 678337 Apr  9 03:34 ceph.log-20220409.gz
-rw-----. 1 ceph ceph 389794 Apr 12 03:09 ceph.log-20220412.gz
-rw-----. 1 ceph ceph 1241977 Apr 13 03:26 ceph.log-20220413.gz
-rw-----. 1 ceph ceph 788096 Apr 14 03:24 ceph.log-20220414.gz
-rw-----. 1 ceph ceph 179467 Apr 20 03:35 ceph.log-20220420.gz
-rw-----. 1 ceph ceph 359751 Apr 21 03:30 ceph.log-20220421.gz
-rw-----. 1 ceph ceph      0 Apr 21 03:31 ceph.log
```

```
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ceph config set global mon_cluster_log_to_journald true
```

```
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ls -ltr ceph.log*
-rw-----. 1 ceph ceph 141884 Apr  8 03:07 ceph.log-20220408.gz
-rw-----. 1 ceph ceph 678337 Apr  9 03:34 ceph.log-20220409.gz
-rw-----. 1 ceph ceph 389794 Apr 12 03:09 ceph.log-20220412.gz
-rw-----. 1 ceph ceph 1241977 Apr 13 03:26 ceph.log-20220413.gz
-rw-----. 1 ceph ceph 788096 Apr 14 03:24 ceph.log-20220414.gz
-rw-----. 1 ceph ceph 179467 Apr 20 03:35 ceph.log-20220420.gz
-rw-----. 1 ceph ceph 359751 Apr 21 03:30 ceph.log-20220421.gz
-rw-----. 1 ceph ceph 2918 Apr 21 19:26 ceph.log
```

#9 - 04/23/2022 09:59 AM - ☐ ☐ ☐

I suspect this issue is due to <https://github.com/ceph/ceph/commit/7c84e06e6f846f6b4b6fd959218b4d474520f429> and have nothing to do with the journald feature.

The above commit keeps an fd to the log file and does not reopen it when logs are rotated. Toggling '`mon_cluster_log_to_journald`' just triggers the '`log_external_close_fds()`' to be called as an side effect, which cause mon to reopen the log file.

Could you try changing some other options like '`mon_cluster_log_to_syslog`', and see if it can also fix the issue?

#10 - 04/25/2022 05:48 PM - Vikhyat Umrao

☐ ☐ wrote:

I suspect this issue is due to <https://github.com/ceph/ceph/commit/7c84e06e6f846f6b4b6fd959218b4d474520f429> and have nothing to do with the journald feature.

Thank you for your feedback. Yes, when I had done a git diff b/w pacific and quincy I was also on these changes but could not figure out that this is the exact fix.

The above commit keeps an fd to the log file and does not reopen it when logs are rotated. Toggling 'mon_cluster_log_to_journald' just triggers the 'log_external_close_fds()' to be called as an side effect, which cause mon to reopen the log file.

Yes, looks like it is exactly the case.

Could you try changing some other options like 'mon_cluster_log_to_syslog', and see if it can also fix the issue?

as you were suspecting the following test confirms it!

```
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ls -ltr ceph.log*
-rw-----. 1 ceph ceph 678337 Apr  9 03:34 ceph.log-20220409.gz
-rw-----. 1 ceph ceph 389794 Apr 12 03:09 ceph.log-20220412.gz
-rw-----. 1 ceph ceph 1241977 Apr 13 03:26 ceph.log-20220413.gz
-rw-----. 1 ceph ceph 788096 Apr 14 03:24 ceph.log-20220414.gz
-rw-----. 1 ceph ceph 179467 Apr 20 03:35 ceph.log-20220420.gz
-rw-----. 1 ceph ceph 359751 Apr 21 03:30 ceph.log-20220421.gz
-rw-----. 1 ceph ceph 361011 Apr 22 03:11 ceph.log-20220422.gz
-rw-----. 1 ceph ceph      0 Apr 22 03:11 ceph.log
```

```
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ceph config get mon mon_cluster_log_to_syslog
default=false
```

```
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ceph config set mon mon_cluster_log_to_syslog false
```

```
[root@gibba001 2c15bd96-7a20-11ec-b7e0-3cecef3d8fb8]# ls -ltr ceph.log*
-rw-----. 1 ceph ceph 678337 Apr  9 03:34 ceph.log-20220409.gz
-rw-----. 1 ceph ceph 389794 Apr 12 03:09 ceph.log-20220412.gz
-rw-----. 1 ceph ceph 1241977 Apr 13 03:26 ceph.log-20220413.gz
-rw-----. 1 ceph ceph 788096 Apr 14 03:24 ceph.log-20220414.gz
-rw-----. 1 ceph ceph 179467 Apr 20 03:35 ceph.log-20220420.gz
-rw-----. 1 ceph ceph 359751 Apr 21 03:30 ceph.log-20220421.gz
-rw-----. 1 ceph ceph 361011 Apr 22 03:11 ceph.log-20220422.gz
-rw-----. 1 ceph ceph 416 Apr 25 17:45 ceph.log
```

Even setting the same default value triggers the call.

#11 - 04/25/2022 05:49 PM - Vikhyat Umrao

If you are okay can you please send a quick fix?

#12 - 04/25/2022 05:49 PM - Vikhyat Umrao

- Status changed from New to Triaged
- Priority changed from Normal to High
- Backport set to quincy

#13 - 04/25/2022 05:51 PM - Vikhyat Umrao

- Project changed from Orchestrator to RADOS
- Category changed from cephadm to Monitor

Looks like it is happening because of mon/LogMonitor changing it back to RADOS.

#14 - 05/02/2022 06:50 PM - Vikhyat Umrao

To me looks like this is the problem?

<https://github.com/ceph/ceph/commit/7c84e06e6f846f6b4b6fd959218b4d474520f429#diff-a61a489ea00d71d720addf06931eefbee07a5ea2eafce6b9e2e321523421b9daR334-R336>

```
if (log_file.empty()) {  
    // do not log this channel  
    fd = -1;  
}
```

#15 - 05/03/2022 09:44 AM -

<https://github.com/ceph/ceph/pull/46124>

Tested locally with

```
kill -SIGHUP <PID of monitor>
```

#16 - 05/03/2022 09:55 AM -

Maybe we should fix the release note(<https://docs.ceph.com/en/latest/releases/quincy/>) first? The workaround there is incorrect.

Should be something like:

<https://tracker.ceph.com/issues/55383> is a known issue - to continue to log cluster log messages to file, run `ceph config set mon mon_cluster_log_to_file true` after every log rotation.

#17 - 05/03/2022 04:27 PM - Vikhyat Umrao

- Status changed from Triaged to Fix Under Review

- Pull request ID set to 46124

#18 - 05/03/2022 04:32 PM - Vikhyat Umrao

wrote:

<https://github.com/ceph/ceph/pull/46124>

Tested locally with

[...]

Thank you.

#19 - 05/03/2022 04:41 PM - Vikhyat Umrao

wrote:

Maybe we should fix the release note(<https://docs.ceph.com/en/latest/releases/quincy/>) first? The workaround there is incorrect.

Should be something like:

<https://tracker.ceph.com/issues/55383> is a known issue - to continue to log cluster log messages to file, run `ceph config set mon mon_cluster_log_to_file true` after every log rotation.

Thanks, completely agree, <https://github.com/ceph/ceph/pull/46128>, done!

#20 - 05/12/2022 05:54 PM - Neha Ojha

- Priority changed from High to Immediate

#21 - 05/23/2022 07:50 PM - Vikhyat Umrao

- Status changed from Fix Under Review to Pending Backport

#22 - 05/23/2022 07:55 PM - Backport Bot

- Copied to Backport #55742: quincy: monitor cluster logs(ceph.log) appear empty until rotated added

#23 - 05/27/2022 03:59 PM - Vikhyat Umrao

- Status changed from Pending Backport to Resolved

#24 - 06/06/2022 06:38 PM - Vikhyat Umrao

Tested with the fixed version and now it is working fine!

```
# date
Thu Jun  2 18:51:49 EDT 2022

# ceph orch upgrade start --image quay.ceph.io/ceph-ci/ceph:quincy
Initiating upgrade to quay.ceph.io/ceph-ci/ceph:quincy

# ceph version
ceph version 17.2.0-361-g95d21f5d (95d21f5dafa91a0443b7910b0b226720b1b86ed6) quincy (stable)

# cd /var/log/ceph/90162cee-77b0-11ec-8f46-44a84203895b/
# ls -ltr ceph.log*
-rw----- 1 ceph ceph  679983 Mar  2 03:33 ceph.log-20220302.gz
-rw----- 1 ceph ceph  262197 Mar  4 03:20 ceph.log-20220304.gz
-rw----- 1 ceph ceph  408546 May  7 03:30 ceph.log-20220507.gz
```

After the upgrade log rotation started working fine!

```
-rw----- 1 ceph ceph  237468 Jun  3 03:33 ceph.log-20220603.gz
-rw----- 1 ceph ceph  664759 Jun  4 03:35 ceph.log-20220604.gz
-rw----- 1 ceph ceph  660498 Jun  5 03:31 ceph.log-20220605.gz
-rw----- 1 ceph ceph  654414 Jun  6 03:18 ceph.log-20220606.gz
-rw----- 1 ceph ceph  4392121 Jun  6 14:35 ceph.log
```