

## ceph-volume - Bug #36492

### ceph lvm list reports wrong json from within a container

10/17/2018 04:02 PM - Sébastien Han

<b>Status:</b>	Resolved	<b>Start date:</b>	10/17/2018
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>		<b>% Done:</b>	0%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>		<b>Reviewed:</b>	
<b>Source:</b>		<b>Affected Versions:</b>	
<b>Tags:</b>		<b>ceph-qa-suite:</b>	
<b>Backport:</b>		<b>Pull request ID:</b>	
<b>Regression:</b>	No		
<b>Severity:</b>	3 - minor		

#### Description

See:

```
[root@magna005 ~]# docker run --rm --privileged --net=host -v /run/lock/lvm:/run/lock/lvm:z -v /dev:/dev -v /etc/ceph:/etc/ceph:z -v /run/lvm/lvmetad.socket:/run/lvm/lvmetad.socket -v /var/lib/ceph/h:/var/lib/ceph/z --entrypoint=ceph-volume 90a28d5c5a38 --cluster ceph lvm list /dev/sda --format=json
stdout: /dev/sdb";";"psC9Rz-ubb0-56ii-3kIK-YJ3t-eLMd-iZUhy6";"vg1";"eM5mmQ-C45O-MeTh-thH5-pyby-qUwr-16shYd
stdout: /dev/sdb";";"psC9Rz-ubb0-56ii-3kIK-YJ3t-eLMd-iZUhy6";"vg1";"A1A1cU-t05h-vL27-E22B-dKiF-qzuj-e9rAC9
stdout: /dev/sdc";";"v1X4gJ-FRZT-e5Gb-0535-gxC1-uHCq-isBBXO";"vg1";"Iz74XQ-E6jJ-GiYW-xZMg-ZYSd-co33-vsCptZ
stdout: /dev/sdc";";"v1X4gJ-FRZT-e5Gb-0535-gxC1-uHCq-isBBXO";"vg1";"A1A1cU-t05h-vL27-E22B-dKiF-qzuj-e9rAC9
stdout: /dev/sdd";";"lbQ9ys-rmee-xs6A-eghN-oy6v-tnYf-YXewn1";"vg1";"Rjq2Kx-PWLX-LPaX-wM1z-de4N-24y4-ZOTkGY
stdout: /dev/sdd";";"lbQ9ys-rmee-xs6A-eghN-oy6v-tnYf-YXewn1";"vg1";"A1A1cU-t05h-vL27-E22B-dKiF-qzuj-e9rAC9
stdout: /dev/sdd";";"lbQ9ys-rmee-xs6A-eghN-oy6v-tnYf-YXewn1";"vg1";"kFmQiy-lC6D-f8Np-feqd-Nkn1-13P6-hPLOWD
stdout: /dev/sdd";";"lbQ9ys-rmee-xs6A-eghN-oy6v-tnYf-YXewn1";"vg1";"xoi1dU-9LED-7goj-0380-2nSF-65Bj-yod5u4
stdout: /dev/sdd";";"lbQ9ys-rmee-xs6A-eghN-oy6v-tnYf-YXewn1";"vg1";"yTDGfg-iujt-qDgq-H18W-r49S-PK58-eYFywM
stdout: /dev/sdd";";"lbQ9ys-rmee-xs6A-eghN-oy6v-tnYf-YXewn1";"vg1";"3HUVo5-0VgM-ZB3v-LW1G-6UTA-vppr-IuIhyr
stdout: /dev/sdd";";"lbQ9ys-rmee-xs6A-eghN-oy6v-tnYf-YXewn1";"vg1";"
stderr: Operation prohibited while global/metadata_read_only is set.
stderr: Recovery of standalone physical volumes failed.
stderr: Cannot process standalone physical volumes
stderr: Operation prohibited while global/metadata_read_only is set.
stderr: Recovery of standalone physical volumes failed.
stderr: Cannot process standalone physical volumes
stderr: Operation prohibited while global/metadata_read_only is set.
stderr: Recovery of standalone physical volumes failed.
stderr: Cannot process standalone physical volumes
stderr: Operation prohibited while global/metadata_read_only is set.
stderr: Recovery of standalone physical volumes failed.
stderr: Cannot process standalone physical volumes
{}

```

No error when listing without a specific block device:

```
[root@magna005 ~]# docker run --rm --privileged --net=host -v /run/lock/lvm:/run/lock/lvm:z -v /de
```

```
v:/dev -v /etc/ceph:/etc/ceph:z -v /run/lvm/lvmetad.socket:/run/lvm/lvmetad.socket -v /var/lib/ceph:/var/lib/ceph/:z --entrypoint=ceph-volume 90a28d5c5a38 --cluster ceph lvm list --format=json {}
```

## History

---

### #1 - 10/17/2018 07:49 PM - Andrew Schoen

The additional logging is coming from lvm and seems to only happen when `ceph-volume lvm list` is run through a container. Perhaps we need to configure lvm differently to allow access this to work through a container. The logs there mention that 'global/metadata\_read\_only is set' but when I look at /etc/lvm/lvm.conf I see 'metadata\_read\_only = 0' which would tell me that it is not set. Does the container not respect lvm.conf?

### #2 - 10/18/2018 12:37 PM - Sébastien Han

'metadata\_read\_only = 0 is set to 0 on the host and in the container. So there must be something that confused lvm inside the container.

### #3 - 10/18/2018 12:43 PM - Sébastien Han

I'd say regardless of the option, can't we just suppress the stdout from the json output? This seems like a bug to me.

### #4 - 10/18/2018 01:10 PM - Alfredo Deza

I agree here that if we are being asked JSON we should return JSON only. The output that is coming from LVM is caused by using the --read-only flag in ceph-volume's internal LVM API calls.

Since we can't remove that flag because it is required for containers, we should suppress this output going to the terminal, and keep it for file logs only (it is a simple flag that should get passed onto process.call

### #5 - 10/18/2018 01:33 PM - Alfredo Deza

We would like to see a Jenkins Job that has this problem to help determine the behavior, and to make sure that our fix can be verified

### #6 - 10/18/2018 08:00 PM - Sébastien Han

Does that mean you cannot start working on fixing this bug unless you have a Jenkins CI job?  
We have an environment that we can use to validate your fix.

I know we have to validate that through CI, however, as you're aware we don't have that kind of job in the CI at the moment. This basically requires 12.2.9 which is not available upstream yet.

Thanks.

### #7 - 10/19/2018 02:39 PM - Alfredo Deza

With the new addition of the development version of containers, can we enable something to verify this works?

### #8 - 10/19/2018 02:42 PM - Sébastien Han

We can build a container image from a branch that would have packages once you push a PR for this.

### #9 - 10/24/2018 02:48 PM - Sébastien Han

```
[root@magna059 ~]# cat /var/log/ceph/ceph-volume.log
[2018-10-24 14:47:04,000][ceph_volume.main][INFO ] Running command: ceph-volume --cluster ceph lvm list --format=json data_vg/data_lv3
[2018-10-24 14:47:04,001][ceph_volume.process][INFO ] Running command: /usr/sbin/lvs --noheadings --readonly --separator=";" -o lv_tags,lv_path,lv_name,vg_name,lv_uuid,lv_size
[2018-10-24 14:47:04,424][ceph_volume.process][INFO ] stdout
```

```
"/dev/data_vg/data_lv1";"data_lv1";"data_vg";"vAil4y-Td4D-ZxXd-2LGM-kyya-voWH-TV3A4h";"600.00g
[2018-10-24 14:47:04,425][ceph_volume.process][INFO ] stdout
"/dev/data_vg/data_lv2";"data_lv2";"data_vg";"c1911y-0awV-fwJE-RUEQ-8s91-tNIV-paJnQi";"600.00g
[2018-10-24 14:47:04,425][ceph_volume.process][INFO ] stdout
"/dev/data_vg/data_lv3";"data_lv3";"data_vg";"2psQvf-bCG4-Bedr-dC1P-gurQ-c99C-6vA5aM";"600.00g
[2018-10-24 14:47:04,425][ceph_volume.process][INFO ] stdout
"/dev/journal_vg/journal_lv1";"journal_lv1";"journal_vg";"MoCerT-Z9v3-YRmY-ukgP-6l6w-IS0R-8VVHLW";"200.00g
[2018-10-24 14:47:04,425][ceph_volume.process][INFO ] stdout
"/dev/journal_vg/journal_lv2";"journal_lv2";"journal_vg";"7twy1m-Yo9i-Dwf8-AOPP-wkt0-h8o8-TkE4YZ";"200.00g
[2018-10-24 14:47:04,425][ceph_volume.process][INFO ] stdout
"/dev/journal_vg/journal_lv3";"journal_lv3";"journal_vg";"1PkngF-bU9V-itDm-LIWq-CaU2-2Y1a-abbYhA";"200.00g
[2018-10-24 14:47:04,425][ceph_volume.process][INFO ] Running command: /usr/sbin/lvs --noheadings --readonly --separator=";" -o
lv_tags,lv_path,lv_name,vg_name,lv_uuid,lv_size
[2018-10-24 14:47:04,830][ceph_volume.process][INFO ] stdout
"/dev/data_vg/data_lv1";"data_lv1";"data_vg";"vAil4y-Td4D-ZxXd-2LGM-kyya-voWH-TV3A4h";"600.00g
[2018-10-24 14:47:04,830][ceph_volume.process][INFO ] stdout
"/dev/data_vg/data_lv2";"data_lv2";"data_vg";"c1911y-0awV-fwJE-RUEQ-8s91-tNIV-paJnQi";"600.00g
[2018-10-24 14:47:04,830][ceph_volume.process][INFO ] stdout
"/dev/data_vg/data_lv3";"data_lv3";"data_vg";"2psQvf-bCG4-Bedr-dC1P-gurQ-c99C-6vA5aM";"600.00g
[2018-10-24 14:47:04,830][ceph_volume.process][INFO ] stdout
"/dev/journal_vg/journal_lv1";"journal_lv1";"journal_vg";"MoCerT-Z9v3-YRmY-ukgP-6l6w-IS0R-8VVHLW";"200.00g
[2018-10-24 14:47:04,831][ceph_volume.process][INFO ] stdout
"/dev/journal_vg/journal_lv2";"journal_lv2";"journal_vg";"7twy1m-Yo9i-Dwf8-AOPP-wkt0-h8o8-TkE4YZ";"200.00g
[2018-10-24 14:47:04,831][ceph_volume.process][INFO ] stdout
"/dev/journal_vg/journal_lv3";"journal_lv3";"journal_vg";"1PkngF-bU9V-itDm-LIWq-CaU2-2Y1a-abbYhA";"200.00g
[2018-10-24 14:47:04,831][ceph_volume.process][INFO ] Running command: /usr/sbin/lvs --noheadings --readonly --separator=";" -o
lv_tags,lv_path,lv_name,vg_name,lv_uuid,lv_size
[2018-10-24 14:47:05,211][ceph_volume.process][INFO ] stdout
"/dev/data_vg/data_lv1";"data_lv1";"data_vg";"vAil4y-Td4D-ZxXd-2LGM-kyya-voWH-TV3A4h";"600.00g
[2018-10-24 14:47:05,211][ceph_volume.process][INFO ] stdout
"/dev/data_vg/data_lv2";"data_lv2";"data_vg";"c1911y-0awV-fwJE-RUEQ-8s91-tNIV-paJnQi";"600.00g
[2018-10-24 14:47:05,211][ceph_volume.process][INFO ] stdout
"/dev/data_vg/data_lv3";"data_lv3";"data_vg";"2psQvf-bCG4-Bedr-dC1P-gurQ-c99C-6vA5aM";"600.00g
[2018-10-24 14:47:05,211][ceph_volume.process][INFO ] stdout
"/dev/journal_vg/journal_lv1";"journal_lv1";"journal_vg";"MoCerT-Z9v3-YRmY-ukgP-6l6w-IS0R-8VVHLW";"200.00g
[2018-10-24 14:47:05,212][ceph_volume.process][INFO ] stdout
"/dev/journal_vg/journal_lv2";"journal_lv2";"journal_vg";"7twy1m-Yo9i-Dwf8-AOPP-wkt0-h8o8-TkE4YZ";"200.00g
[2018-10-24 14:47:05,212][ceph_volume.process][INFO ] stdout
"/dev/journal_vg/journal_lv3";"journal_lv3";"journal_vg";"1PkngF-bU9V-itDm-LIWq-CaU2-2Y1a-abbYhA";"200.00g
[2018-10-24 14:47:05,212][ceph_volume.process][INFO ] Running command: /usr/sbin/pvs --no-heading --readonly --separator=";" -o
pv_name,pv_tags,pv_uuid,vg_name,lv_uuid
[2018-10-24 14:47:05,618][ceph_volume.process][INFO ] stdout
/dev/sdb1";"1zhkMS-6zSD-Dwfr-Bkcr-pnd5-nOf-Xk0GRG";"data_vg";"vAil4y-Td4D-ZxXd-2LGM-kyya-voWH-TV3A4h
[2018-10-24 14:47:05,618][ceph_volume.process][INFO ] stdout
/dev/sdb1";"1zhkMS-6zSD-Dwfr-Bkcr-pnd5-nOf-Xk0GRG";"data_vg";"2psQvf-bCG4-Bedr-dC1P-gurQ-c99C-6vA5aM
[2018-10-24 14:47:05,619][ceph_volume.process][INFO ] stdout
/dev/sdc1";"HGniou-f9Gl-kbPX-122U-8Xzg-WDWL-Z0v4Wn";"data_vg";"c1911y-0awV-fwJE-RUEQ-8s91-tNIV-paJnQi
[2018-10-24 14:47:05,619][ceph_volume.process][INFO ] stdout
/dev/sdc1";"HGniou-f9Gl-kbPX-122U-8Xzg-WDWL-Z0v4Wn";"data_vg";"2psQvf-bCG4-Bedr-dC1P-gurQ-c99C-6vA5aM
[2018-10-24 14:47:05,619][ceph_volume.process][INFO ] stdout /dev/sdc1";"HGniou-f9Gl-kbPX-122U-8Xzg-WDWL-Z0v4Wn";"data_vg";"
[2018-10-24 14:47:05,619][ceph_volume.process][INFO ] stdout
/dev/sdd1";"wHN7at-dBji-1MwE-sRRs-OcOB-kRJ9-yLDfCP";"journal_vg";"MoCerT-Z9v3-YRmY-ukgP-6l6w-IS0R-8VVHLW
[2018-10-24 14:47:05,619][ceph_volume.process][INFO ] stdout
/dev/sdd1";"wHN7at-dBji-1MwE-sRRs-OcOB-kRJ9-yLDfCP";"journal_vg";"7twy1m-Yo9i-Dwf8-AOPP-wkt0-h8o8-TkE4YZ
[2018-10-24 14:47:05,619][ceph_volume.process][INFO ] stdout
/dev/sdd1";"wHN7at-dBji-1MwE-sRRs-OcOB-kRJ9-yLDfCP";"journal_vg";"1PkngF-bU9V-itDm-LIWq-CaU2-2Y1a-abbYhA
[2018-10-24 14:47:05,619][ceph_volume.process][INFO ] stdout /dev/sdd1";"wHN7at-dBji-1MwE-sRRs-OcOB-kRJ9-yLDfCP";"journal_vg";"
[2018-10-24 14:47:05,619][ceph_volume.process][INFO ] stderr Operation prohibited while global/metadata_read_only is set.
[2018-10-24 14:47:05,619][ceph_volume.process][INFO ] stderr Recovery of standalone physical volumes failed.
[2018-10-24 14:47:05,619][ceph_volume.process][INFO ] stderr Cannot process standalone physical volumes
[2018-10-24 14:47:05,620][ceph_volume.process][INFO ] stderr Operation prohibited while global/metadata_read_only is set.
[2018-10-24 14:47:05,620][ceph_volume.process][INFO ] stderr Recovery of standalone physical volumes failed.
[2018-10-24 14:47:05,620][ceph_volume.process][INFO ] stderr Cannot process standalone physical volumes
[2018-10-24 14:47:05,620][ceph_volume.process][INFO ] stderr Operation prohibited while global/metadata_read_only is set.
[2018-10-24 14:47:05,620][ceph_volume.process][INFO ] stderr Recovery of standalone physical volumes failed.
[2018-10-24 14:47:05,620][ceph_volume.process][INFO ] stderr Cannot process standalone physical volumes
[2018-10-24 14:47:05,620][ceph_volume.devices.lvm.listing][WARNING] device is not part of ceph: data_vg/data_lv3
```

**#10 - 10/24/2018 02:58 PM - Alfredo Deza**

- Status changed from New to Need Review

master PR <https://github.com/ceph/ceph/pull/24738>

**#11 - 10/25/2018 01:28 PM - Alfredo Deza**

- Status changed from Need Review to Resolved

mimic PR <https://github.com/ceph/ceph/pull/24740>

luminous PR <https://github.com/ceph/ceph/pull/24741>

**#12 - 12/05/2018 01:08 AM - Francois Lafont**

Hi @all. :)

I'm trying to deploy a Ceph Mimic cluster in docker containers via ceph-ansible and I'm currently stuck with this bug exactly.

Is there a workaround or should I wait the next Mimic release?

**#13 - 12/10/2018 11:57 AM - Alfredo Deza**

Correct, there hasn't been a Mimic release that includes the fixes here. You will have to wait for one, hopefully soon!

**#14 - 12/10/2018 01:03 PM - Francois Lafont**

Ok, no problem.

Thx for your answer Alfredo. For the moment, I have built "my" Mimic ceph/daemon docker container with the little fix and it works great with ceph-ansible. :)