

OpenNebula - Bug #3726

Ceph monitor check broken during normal cluster degradation

03/30/2015 02:34 PM - Joel Merrick

Status:	Closed	Start date:	03/30/2015
Priority:	High	Due date:	
Assignee:	Javi Fontan	% Done:	0%
Category:	Documentation	Estimated time:	0.00 hour
Target version:	Release 4.12.1	Pull request:	
Resolution:			
Affected Versions:	OpenNebula 4.12		

Description

There is a bug in the new Ceph monitor. It's assuming the columns are always static in the output of the ceph/rados output are static - however when there's any degradation (perfectly normal if you've lost a disk and rebalancing happening) then the awk picks up the wrong column (0 in my case). See <https://gist.github.com/joelio/64ae2b9fe9116fcca4c6>

This is a serious fault as it stops all provisioning on an otherwise healthy cluster.

I think it'd be more prudent to use the JSON output and parse that properly?

History

#1 - 03/30/2015 02:38 PM - Joel Merrick

Been notified that XML output supported

```
ceph df -f xml
```

Might be better than JSON for ONE usage

#2 - 03/30/2015 02:52 PM - Ruben S. Montero

Are you running 0.87.1?

Anyway, using xml/json should be better

#3 - 03/30/2015 02:53 PM - Ruben S. Montero

- Target version set to Release 4.14

#4 - 03/30/2015 03:03 PM - Joel Merrick

Yea, current stable (0.87 Giant)

ceph:

Installed: 0.87.1-1trusty

Candidate: 0.87.1-1trusty

Version table: *** 0.87.1-1trusty 0

999 <http://ceph.com/debian-giant/> trusty/main amd64 Packages

#5 - 03/30/2015 03:11 PM - Ruben S. Montero

- Status changed from Pending to New
- Target version changed from Release 4.14 to Release 4.12.1

#6 - 03/30/2015 03:13 PM - Ruben S. Montero

- Assignee set to Javi Fontan

#7 - 03/30/2015 04:05 PM - Ruben S. Montero

Hi,

Checking it; I have

```
# rados df
pool name      category      KB  objects  clones  degraded  unfound  rd  rd KB  wr  wr KB
...
```

but with ceph:

```
# ceph df
GLOBAL:
  SIZE  AVAIL  RAW USED  %RAW USED
...
POOLS:
  NAME      ID  USED  %USED  MAX AVAIL  OBJECTS
...
```

so you say that ceph df becomes rados df under degradation?

Cheers

#8 - 03/30/2015 04:23 PM - Joel Merrick

I'm noticing no MAX_AVAIL in my output too

```
root@vm-head-01:/var/log/one# ceph df
GLOBAL:
  SIZE  AVAIL  RAW USED  %RAW USED
59304G  53306G  5960G    10.05
POOLS:
  NAME      ID  USED  %USED  MAX AVAIL  OBJECTS
rbd         0   0    0     0         0
cephfs_data 1 1930G  3.26   0 2062873
cephfs_metadata 2 38567k 0      0 75344
one        3 44924M 0.07   0 11350
```

#9 - 03/30/2015 04:24 PM - Joel Merrick

root@vm-head-01:/var/log/one# ceph df

GLOBAL:

SIZE	AVAIL	RAW USED	%RAW USED
59304G	53306G	5960G	10.05

POOLS:

NAME	ID	USED	%USED	MAX AVAIL	OBJECTS
rbd	0	0	0	0	0
cephfs_data	1	1930G	3.26	0	2062873
cephfs_metadata	2	38567k	0	0	75344
one	3	44924M	0.07	0	11350

#10 - 03/30/2015 04:26 PM - Joel Merrick

This is a fresh install too, nothing further done with Ceph for opennebula apart from setting auth up.

If it's just a case of getting MAX_AVAIL working then need to understand what we need to do and add to the docs :)

#11 - 03/30/2015 04:48 PM - Joel Merrick

Ok, so I'm thinking this is a bug in Ceph now, sorry for the confusion!

<http://tracker.ceph.com/issues/10257>

#12 - 03/30/2015 04:57 PM - Joel Merrick

Yes, I can confirm this bug with Ceph, not ONE.

You can close (or perhaps mark documentation?) this bug as there's not a lot you guys can do to mitigate.

Basically ensure ALL OSDs are in and up or do not use 0.87.1. specifically - looks fixed in later releases.

Thanks again guys! :)

#13 - 03/30/2015 06:45 PM - Ruben S. Montero

- *Category changed from Drivers - Storage to Documentation*

Great! I'm no 100% sure but I think this bug also appears when the OSD's are weight'ed to 0... A note to the docs would be nice, at least in the release note pages.

Thanks for letting us know

Cheers

Joel Merrick wrote:

Yes, I can confirm this bug with Ceph, not ONE.

You can close (or perhaps mark documentation?) this bug as there's not a lot you guys can do to mitigate.

Basically ensure ALL OSDs are in and up or do not use 0.87.1. specifically - looks fixed in later releases.

Thanks again guys! :)

#14 - 04/27/2015 03:26 PM - Ruben S. Montero

- Status changed from New to Closed
- Resolution set to worksforme

#15 - 04/27/2015 03:26 PM - Ruben S. Montero

- Status changed from Closed to Pending
- Resolution deleted (worksforme)

#16 - 05/05/2015 03:57 PM - Jaime Melis

- Status changed from Pending to Closed

Documented