

OpenNebula - Bug #4278

Wrong OpenNebula DB consistency regarding to volatile disk quotas

01/07/2016 02:24 PM - Esteban Freire Garcia

Status:	Closed	Start date:	01/07/2016
Priority:	Normal	Due date:	
Assignee:	Ruben S. Montero	% Done:	0%
Category:	Core & System	Estimated time:	0.00 hour
Target version:	Release 5.0	Pull request:	
Resolution:	fixed		
Affected Versions:	OpenNebula 4.12		

Description

- We are observing wrong volatile disks quotas from time to time and since some months ago.

- At the beginning, we thought it could be due to some crash on the OpenNebula head-node or some kind of connectivity lost with the OpenNebula database. But after being observing this behaviour for a while now, we can confirm that it is not due to any crash at OpenNebula head-node or any connectivity lost with the OpenNebula database.

- Some examples:

- We executed the onedb fsck command yesterday, this was the output:

```
...
onedb fsck [ ... ]
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql_2016-1-6_9:4:43.bck
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file
```

```
User 946 quotas: VOLATILE_SIZE_USED has 1024 is 0
User 946 quotas: Image 1781 RVMS has 0 is 1
User 951 quotas: VOLATILE_SIZE_USED has 20480 is 0
User 1207 quotas: VOLATILE_SIZE_USED has 10240 is 0
User 1216 quotas: VOLATILE_SIZE_USED has 10240 is 0
Group 411 quotas: VOLATILE_SIZE_USED has 1024 is 0
Group 411 quotas: Image 1781 RVMS has 0 is 1
Group 420 quotas: VOLATILE_SIZE_USED has 20480 is 0
Group 463 quotas: Datastore 104 IMAGES_USED has -2 is 0
Group 463 quotas: Datastore 104 SIZE_USED has -950 is 0
Group 465 quotas: VOLATILE_SIZE_USED has 10240 is 0
Group 468 quotas: VOLATILE_SIZE_USED has 10240 is 0
```

Total errors found: 12

A copy of this output was stored in /var/log/one/onedb-fsck.log

```
onedb fsck [ ... ]
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql_2016-1-6_9:6:0.bck
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file
```

Total errors found: 0

A copy of this output was stored in /var/log/one/onedb-fsck.log

...

- According to previous output, it fixed quotas information for group 465 and user 1207 but it was not like that and we needed to execute the onedb fsck command today at first hour again:

```
...
onedb fsck [ ... ]
MySQL dump stored in /var/lib/one/mysql_localhost_opennebula.sql_2016-1-7_9:0:12.bck
Use 'onedb restore' or restore the DB using the mysql command:
mysql -u user -h server -P port db_name < backup_file
```

```
User 794 quotas: VOLATILE_SIZE_USED has 10240 is 0
User 1207 quotas: VOLATILE_SIZE_USED has 71680 is 0
User 1207 quotas: Image 1769 RVMS has 0 is 1
User 1207 quotas: Image 1908 RVMS has 0 is 1
User 1207 quotas: Datastore 106 IMAGES_USED has 1 is 2
User 1207 quotas: Datastore 106 SIZE_USED has 2048 is 62248
Group 331 quotas: VOLATILE_SIZE_USED has 10240 is 0
Group 465 quotas: VOLATILE_SIZE_USED has 71680 is 0
Group 465 quotas: Image 1769 RVMS has 0 is 1
Group 465 quotas: Image 1908 RVMS has 0 is 1
Group 465 quotas: Datastore 106 IMAGES_USED has 2 is 3
Group 465 quotas: Datastore 106 SIZE_USED has 12048 is 72248
```

```
Total errors found: 12
A copy of this output was stored in /var/log/one/onedb-fsck.log
```

```
...
- As it can be seen, group 465 still had same yesterday issue which it was supposed to be solved. After this, these groups are showing the right volatile disk quota. Some examples:
```

```
...
[root@opennebula4 ~]# onegroup show 331
GROUP 331 INFORMATION
ID      : 331
NAME    :

GROUP TEMPLATE
GROUP_ADMIN_DEFAULT_VIEW="groupadmin"
GROUP_ADMIN_VIEWS="groupadmin"
SUNSTONE_VIEWS="cloud,user"
```

```
USER ID ADMIN
794 *
```

```
RESOURCE USAGE & QUOTAS
NUMBER OF VMS      MEMORY      CPU      VOLATILE_SIZE
2 / - 32.3G / - 100.00 / - 0M / 10G
```

```
DATASTORE ID      IMAGES      SIZE
104 2 / - 1020M / 15G
106 0 / - 0M / 5G
```

```
NETWORK ID      LEASES
0 2 / -

IMAGE ID      RUNNING VMS
```

```
557 1 / -
1073 1 / -
...
...

```

```
[root@opennebula4 ~]# onegroup show 465
GROUP 465 INFORMATION
ID      : 465
NAME    :
```

```
GROUP TEMPLATE
GROUP_ADMIN_DEFAULT_VIEW="user"
GROUP_ADMIN_VIEWS="groupadmin"
SUNSTONE_VIEWS="cloud,user"
```

```
USER ID ADMIN
1207 *
1208 *
```

```
RESOURCE USAGE & QUOTAS
NUMBER OF VMS      MEMORY      CPU      VOLATILE_SIZE
1 / - 950G / - 40.00 / - 0M / 100G
```

```
DATASTORE ID      IMAGES      SIZE
119 1 / - 510M / 40G
106 3 / - 70.6G / 2T
```

```
NETWORK ID      LEASES
0 1 / -
```

```
IMAGE ID      RUNNING VMS
1785 1 / -
1769 1 / -
1908 1 / -
...

```

- But we did not have any crash or connectivity lost since we executed onedb fsck command yesterday and we executed it again today so it is because we are considering it as a bug.

Thanks in advance,
Esterban

Associated revisions

Revision 87709d93 - 06/14/2016 09:23 AM - Ruben S. Montero

bug #4278: Sync volatile disk computations in core & fsck

History

#1 - 02/12/2016 11:01 AM - Anonymous

We are still seeing this with release 4.14.
The result is that we need to fsck the database 2 times a week due to inconsistent states.

Esteban told me this issue is scheduled to be fixed in version 5, can i get any confirmation on that (to prevent it being forgotten)

#2 - 02/18/2016 01:25 AM - Ruben S. Montero

- *Target version set to Release 5.0*

Target version is 5.0 in redmine, now. We'll look at it thanks!

#3 - 02/18/2016 08:40 AM - Esteban Freire Garcia

Thank you very much Ruben :)

#4 - 05/17/2016 02:24 PM - Ruben S. Montero

- *Assignee set to Carlos Martín*

#5 - 06/09/2016 09:26 AM - Ignacio M. Llorente

- *Assignee changed from Carlos Martín to Ruben S. Montero*

#6 - 06/14/2016 09:33 AM - Ruben S. Montero

- *Status changed from Pending to Closed*

- *Resolution set to fixed*

The volatile disk size was slightly different in core and fsck. This could impact on DS cloning on it self, like Ceph... I think this can be the problem. I am closing this and if this patch does not solve it we'll reopen