

OpenNebula - Feature #4885

Ceph Snapshots cannot be deleted

10/29/2016 09:31 PM - Roy Keene

Status:	Closed	Start date:	10/29/2016
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:	Drivers - Storage	Estimated time:	0.00 hour
Target version:	Release 5.4	Pull request:	
Resolution:	fixed		
Description			
<p>In OpenNebula 5.2.0, disk snapshots for disks in Ceph cannot be deleted.</p> <p>You can't delete the most recent snapshot since it is ACTIVE, and you can't delete any other snapshots because a tree is formed and it has children.</p> <p>This whole concept does not match Ceph's snapshots at all, there is no "active" snapshot (snapshots are read-only) and you can delete any snapshot without regards to order.</p> <p>Attached is a simple patch which removes this notion altogether and fixes snapshots for Ceph (as can be seen here: http://www.rkeene.org/tmp/opennebula-ceph-snapshots.mp4)</p> <p>I only use Ceph for storage so the patch may not be generally useful in its current state.</p>			
Related issues:			
Duplicates Feature # 4008: option to unprotect the delete action on disk snap...		Closed	09/29/2015

Associated revisions

Revision ebc165f7 - 07/05/2017 04:13 PM - Ruben S. Montero

F #4885: Enable TM mads to define orphan snapshots. This is control under TM_MAD_CONF for each MAD. By default orphans are not allowed

History

#1 - 10/30/2016 11:07 PM - Ruben S. Montero

- Tracker changed from Bug to Feature
- Status changed from Pending to New
- Target version set to Release 5.4

Hi Roy

Agree, preserve snapshots with children maybe too restrictive for some backends. During the OpenNebulaConf last week we've discussed this with the StorPool team, because their storage system does not have this limitation either.

Our conclusion was to add a configuration attribute for the driver (e.g. TM_MAD=[NAME="ceph", ALLOW_ORPHANS="yes"...] or similar) in oned.conf. This is basically adding an if for the current behavior and the one you propose.

Thanks

PS: Moved to the feature tracker

#2 - 01/08/2017 10:44 AM - Bernhard J. M. Grün

Hey,

it seems this feature request is somehow linked to that feature request #4008.

I would really like to see this solved too.

Regards

#3 - 03/28/2017 12:48 AM - kvaps kvaps

Hello, curenly I pending new storage driver based on qcow2, and I have same problem.

I know the way to simple live-squashing of disk-snapshots:

https://github.com/kvaps/opennebula-addon-raw_qcow2/blob/master/tm/raw_qcow2/snap_delete#L74

I will correct qcow2 driver too after, so I that this is wrong limitation, we need to remove it altogether.

#4 - 04/07/2017 03:08 PM - Gerben Meijer

This decoupling of parent/child relationship is needed for our ScaleIO backend as well.

#5 - 04/27/2017 08:54 AM - Roman Bogachev

Ruben S. Montero wrote:

Our conclusion was to add a configuration attribute for the driver (e.g. TM_MAD=[NAME="ceph", ALLOW_ORPHANS="yes"...] or similar) in oned.conf. This is basically adding an if for the current behavior and the one you propose.

Did I understand right, that can use the next configuration in oned.conf?

e.g.

```
TM_MAD = [  
    EXECUTABLE = "one_tm",  
    ARGUMENTS = "-t 15 -d dummy,lvm,shared,fs_lvm,qcow2,ssh,ceph,dev,vcenter,iscsi_libvirt"  
    ALLOW_ORPHANS="yes"  
]
```

#6 - 04/27/2017 02:21 PM - Gerben Meijer

Roman Bogachev wrote:

Did I understand right, that can use the next configuration in oned.conf?

No, that is the intention; there is no code nor patch yet that does that.

#7 - 07/05/2017 04:13 PM - Ruben S. Montero

- *Duplicates Feature #4008: option to unprotect the delete action on disk snapshots added*

#8 - 07/05/2017 04:14 PM - Ruben S. Montero

- *Status changed from New to Closed*

- *Resolution set to fixed*

This is now implemented for 5.4

Those developing drivers for storage backends, need to add an attribute ALLOW_ORPHANS="YES" in the TM_MAD_CONF stanza for the storage. By default, OpenNebula will set Ceph to allow orphans.

In order to not introduce confusion, current running VMs will preserve the semantics used when the snapshots were taken.

Files

opennebula-5.2.0-fixdisksnapshots.diff	914 Bytes	10/29/2016	Roy Keene
--	-----------	------------	-----------