

## OpenNebula - Bug #3937

### After disk snapshot with suspend / resume vlan does not get (re-)tagged on openvswitch

08/14/2015 03:43 PM - Stefan Kooman

<b>Status:</b>	Closed	<b>Start date:</b>	08/14/2015
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assignee:</b>		<b>% Done:</b>	0%
<b>Category:</b>	Drivers - VM	<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	Release 4.14	<b>Pull request:</b>	
<b>Resolution:</b>	fixed		
<b>Affected Versions:</b>	OpenNebula 4.12		
<b>Description</b>			
<p>The the vlan does not get (re-)tagged after the VM gets resumed again (after a snapshot create action for a VM (with a RAW image):</p> <p>ovs-vsctl show:</p> <pre>Port "vnet1"     Interface "vnet1"</pre> <p>While it was</p> <pre>Port "vnet1"     tag: 228     Interface "vnet1"</pre> <p>before the "DISK_SNAPSHOT" / suspend action</p>			

#### Associated revisions

##### Revision 988cf671 - 09/01/2015 09:32 AM - Jaime Melis

Bug #3937: After disk snapshot with suspend / resume vlan does not get (re-)tagged on openvswitch

##### Revision 40caf75b - 09/02/2015 01:35 PM - Jaime Melis

Bug #3937: Apply network drivers after disk-snapshot-revert

#### History

##### #1 - 08/21/2015 04:01 PM - Javi Fontan

- Category set to Drivers - VM
- Target version set to Release 4.14

##### #2 - 09/01/2015 09:40 AM - Jaime Melis

- Status changed from Pending to Closed
- Resolution set to fixed
- Affected Versions OpenNebula 4.12 added

### #3 - 09/01/2015 03:39 PM - Stefan Kooman

I replaced `/var/lib/one/remotes/vmm/one_vmm_exec.rb` with this new version and did a "onehost sync --force" afterwards. Tag does not get re-applied. Besides that, it *looks* like the VM does not get a "poweroff --hard" before the revert ... leading to OS crash. Is this the right way to test this fix?

### #4 - 09/02/2015 09:07 AM - Stefan Kooman

I just checked out master and recompiled / reinstalled (`/usr/lib/one/mads/one_vmm_exec.rb` is also the new version). The VM is shut down but seems to be resumed as there is no normal boot sequence when the VM is running again (bios -> boot). After a reboot the VM ends in a strack trace with ext4 inode errors ...

### #5 - 09/02/2015 02:14 PM - Jaime Melis

- Status changed from Closed to Assigned

I have updated added a part that was missing, in order to re-apply also on revert, not only on create. It should not be necessary to run onehost sync or to reinstall, just replace the `one_vmm_exec.rb` and restart opennebula.

I don't understand exactly what you are doing. My workflow is as follows:

- VM is running
- `onevm disk-snapshot-create <vmid> <diskid> <snapshot name>`
- I observe how the VM disappears from libvirt for a second as it's being suspended, and then reappears and the vnm drivers are reapplied (testing with ovschance)
- `onevm disk-snapshot-revert ...`
- I observe the same thing as with `disk-snapshot-create`

Note that there is no `poweroff --hard` involved here, I'm doing this while the VM is running, and after the operation the VM is running again. OpenNebula does the suspend behind the scenes, I only need to instruct it to do `disk-snapshot-create`.

Can you clarify what you mean with your previous comments? Maybe posting your workflow will help.

### #6 - 09/02/2015 03:10 PM - Stefan Kooman

Note that there is no `poweroff --hard` involved here, I'm doing this while the VM is running, and after the operation the VM is running again. OpenNebula does the suspend behind the scenes, I only need to instruct it to do `disk-snapshot-create`.

I think that's the problem: if you replace the root disk of a VM (rootfs) with a previous snapshot and resume the VM again, you will end up with a corrupted filesystem. The system expects files / inodes / fscache at certain places ... and all of the sudden it's gone or somewhere else. I believe the correct way to do a "snapshot\_revert\_while\_running" is to do a "poweroff --hard" -> `onevm disk-snapshot-revert` -> `poweron`.

### #7 - 09/02/2015 04:33 PM - Stefan Kooman

TL;DR The vlan tags get applied nicely, so the bug is fixed.

I assumed a "poweroff --hard" would have been given ... to avoid fs corruption to the running fs. But apparently it's not how it's designed. Maybe a warning should be added that a revert for mounted filesystems is very dangerous and will lead to dataloss.

**#8 - 09/03/2015 11:37 AM - Jaime Melis**  
*- Status changed from Assigned to Closed*