

OpenNebula - Feature #3782

Ceph snapshots

04/29/2015 07:50 PM - Karsten Nielsen

| | | | |
|---|------------------|------------------------|------------|
| Status: | Closed | Start date: | 04/29/2015 |
| Priority: | Sponsored | Due date: | |
| Assignee: | Ruben S. Montero | % Done: | 0% |
| Category: | Core & System | Estimated time: | 0.00 hour |
| Target version: | Release 4.14 | Pull request: | |
| Resolution: | fixed | | |
| Description | | | |
| <p>We are using ceph as our storage platform for all that we do.</p> <p>We would like to be able to snapshot ceph through opennebula.</p> <p>First usecase can be described as follows:</p> <p>We use the opennebula api to deploy a VM. We define that the VM should have additional rbd's attached as storage devices for docker running inside the VM.</p> <p>After a while we use salt-stack to upgrade the docker image running on the additional rbd storage devices. Before we upgrade the docker image we create a snapshot of the rbd device through the opennebula api. After the upgrade we find some problem with the upgraded docker image and run salt-stack to revert the rbd to the previous snapshot of the rbd through the opennebula api.</p> <p>After some time of operation of the VM I login to the sunstone interface and browse to the VM and browse the storage attached to the where I can see the number of snapshots and the amount of space that it occupies.</p> <p>Second usecase:</p> <p>We operate a regular VM that we need to do some major upgrades to, we start by doing a snapshot of the VM's storage, either through sunstone or opennebula api. Do the upgrade and try to verify that the upgrade went as expected.</p> <p>The upgrade failed and we revert to the snapshot before the upgrade, either through sunstone or opennebula api.</p> | | | |

Associated revisions

Revision 676bfcc1 - 05/14/2015 05:43 PM - Ruben S. Montero

feature #3782: Initial class for Snapshot list

Revision 98f76cf3 - 05/19/2015 04:41 PM - Ruben S. Montero

feature #3782: Create and Revert operations for Disk Snapshots. Core and OCA levels

Revision e3066807 - 05/19/2015 10:21 PM - Ruben S. Montero

feature #3782: Add DISK_SNAPSHOT_POWEROFF state

Revision 986cf14c - 05/20/2015 03:18 PM - Javi Fontan

feature #3782: add snapshots to onevm show

Revision 20fac744 - 05/20/2015 03:48 PM - Ruben S. Montero

feature #3782: Generate snapshot files for TransferManager

Revision 904b3248 - 05/21/2015 12:12 PM - Ruben S. Montero

feature #3782: Remove SOURCE and PARENT from snapshots. Drivers MUST generate internal snapshot references from snapshot_id, parent_id

Revision 9bb5770f - 05/21/2015 12:52 PM - Ruben S. Montero

feature #3782: Snapshot event triggers

Revision a5c04b47 - 05/26/2015 09:24 AM - Ruben S. Montero

feature #3782: Re-design of revert operation. New action for drivers, and new state.

Revision 4f473093 - 05/26/2015 10:52 AM - Ruben S. Montero

feature #3782: Deal with snapshot id=0

Revision 036862b1 - 05/26/2015 11:52 AM - Ruben S. Montero

feature #3782: Add children list. Rename PARENT_ID to PARENT

Revision 83ac6188 - 05/27/2015 02:29 PM - Ruben S. Montero

feature #3782: New state to delete snapshots. Delete action OCA (Ruby), core and CLI

Revision b0f72d5a - 05/28/2015 03:19 PM - Jaime Melis

Feature #3782: Add ceph disk snapshot actions

Revision f3c70ff3 - 05/28/2015 03:30 PM - Jaime Melis

Feature #3782: remove debug info

Revision 4929b50b - 05/29/2015 12:02 PM - Jaime Melis

Feature #3782: add support for snapshots of persistent images in ceph

Revision 3f54b6f6 - 05/30/2015 11:29 AM - Ruben S. Montero

feature #3782: Update copy constructor for Template. It should purge

before copying. Adds a generic clear() method, and use it in class implementation

Revision c2bf13fe - 05/30/2015 11:31 AM - Ruben S. Montero

feature #3782: Copy constructor for Snapshots. Used to transfer the snapshot link from VM disk to Image objects

Revision 9e30da08 - 05/30/2015 11:33 AM - Ruben S. Montero

feature #3782: Add snapshots to Image object. Transfer them back from VM disk

Revision 82eac63f - 06/01/2015 05:25 PM - Ruben S. Montero

feature #3782: Assignment copy constructor for Snapshots. Clear disk_id

Revision a841b12e - 06/01/2015 05:26 PM - Ruben S. Montero

feature #3782: Copy assignment operator for Templates

Revision e599187a - 06/01/2015 05:28 PM - Ruben S. Montero

feature #3782: Copy snapshots from Image to VM disks, and back. Still needs to do the same on attach detach operations

Revision da794476 - 06/01/2015 11:50 PM - Ruben S. Montero

feature #3782: Transfer snapshot list from/to attach disks

Revision 14da3131 - 06/02/2015 06:45 PM - Ruben S. Montero

feature #3782: Image with snapshots cannot be made non-persistent

Revision 43be9c66 - 06/02/2015 09:12 PM - Ruben S. Montero

feature #3782: Image with snapshots can be only of type OS or DATABLOCK

Revision 023c853b - 06/02/2015 09:42 PM - Ruben S. Montero

feature #3782: Cannot clone images with snapshots

Revision 8978f56c - 06/03/2015 12:37 PM - Jaime Melis

feature #3782: Display snapshots in oneimage show

Revision 5510b56b - 06/04/2015 06:17 PM - Ruben S. Montero

feature #3782: Delete operations are now two-phase. More checks on disk-snapshots actions. Improve const-ness of Snapshot class

Revision 480ad917 - 06/04/2015 06:17 PM - Ruben S. Montero

feature #3782: Add missing unlocks to failure exits

Revision 28d3dcb8 - 06/05/2015 01:15 AM - Ruben S. Montero

feature #3782: Snapshot operations for Image class. Add target_snapshot to Image class to identify operation on progress

Revision 1bd6ffe6 - 06/05/2015 01:18 AM - Ruben S. Montero

feature #3782: Delete operations for image snapshots. Core image manager and driver functions

Revision 62a98c55 - 06/05/2015 01:24 AM - Ruben S. Montero

feature #3782: Image XML-RPC methods for snapshot delete

Revision 78c7bdbd - 06/05/2015 01:27 AM - Ruben S. Montero

feature #3782: Ruby OCA method for snapshot-delete, cli command. Datastore driver action for snapshot delete

Revision 9348a0d4 - 06/05/2015 02:45 AM - Ruben S. Montero

feature #3782: Revert action for images

Revision 0bbb99a2 - 06/05/2015 11:32 AM - Ruben S. Montero

feature #3782: Image are in LOCKED state while performing snapshot operations

Revision 08356ed8 - 06/05/2015 02:01 PM - Ruben S. Montero

feature #3782: Flatten action for Images

Revision 211619d8 - 06/08/2015 10:04 AM - Jaime Melis

Feature #3782: Action for Ceph snapshot flatten

Revision c3621e33 - 06/08/2015 10:45 AM - Jaime Melis

Feature #3782: Update install.sh and remotes with ds snap actions

Revision 58a260f9 - 06/08/2015 01:00 PM - Jaime Melis

Feature #3782: Add dummy snapshot_* placeholders for all the TM drivers but Ceph

Revision ee1fb2c6 - 06/08/2015 05:59 PM - Ruben S. Montero

feature #3782: Do not snapshot volatile disks

Revision c3741e4f - 06/08/2015 10:21 PM - Ruben S. Montero

feature #3782: Remove snapshot list if empty

Revision 9d20167c - 06/09/2015 09:25 PM - Ruben S. Montero

feature #3782: Remove SAVE_AS from core

Revision ce851880 - 06/09/2015 11:30 PM - Ruben S. Montero

feature #3782: Changes in OCA (Ruby) and CLI. Updated direct references to OCA methods

Revision 630e0360 - 06/10/2015 10:53 AM - Ruben S. Montero

feature #3782: Removed _hot from saveas functions

Revision 80c06dca - 06/10/2015 05:56 PM - Ruben S. Montero

feature #3782: renamed method to disksaveas, cli and Ruby OCA. Added snap id to save as disks

Revision 943e0e52 - 06/11/2015 01:24 PM - Javi Fontan

feature #3782: add qcow2 tm/snap_{create,revert}

Revision e5bcaea1 - 06/11/2015 01:24 PM - Javi Fontan

feature #3782: add qcow2 tm/snap_delete

Revision 996ee89d - 06/11/2015 04:12 PM - Jaime Melis

feature #3782: allow disk-saveas operations on the currently running disk

Revision ae33c915 - 06/11/2015 04:12 PM - Jaime Melis

feature #3782: Add snapshot id to the onevm disk-saveas command

Revision d8c15a13 - 06/11/2015 04:12 PM - Jaime Melis

feature #3782: Update ceph cpds and mvds scripts to handle snaps

Revision 8ec736d5 - 06/11/2015 04:22 PM - Jaime Melis

feature #3782: Update all cpds to include the SNAP_ID parameter

Revision b5136a15 - 06/12/2015 03:50 PM - Jaime Melis

Feature #3782: Remove debugging data

Revision 1b531f82 - 06/12/2015 03:50 PM - Jaime Melis

Feature #3782: Simplify all TM/mvds action as now they don't handle deferred saveas

Revision 4a0438c1 - 06/15/2015 09:28 PM - Ruben S. Montero

feature #3782: Fix comment

Revision 32832269 - 06/15/2015 11:59 PM - Ruben S. Montero

feature #3782: Fix return codes

Revision dac9db95 - 06/16/2015 09:59 PM - Ruben S. Montero

feature #3782: Size for snapshots (Virtual size). Renamed is* functions to is_* in VirtualMachine.

Revision 9542d396 - 06/17/2015 08:10 AM - Ruben S. Montero

feature #3782: Update quotas on non-persistent snapshots

Revision b2101281 - 06/17/2015 02:02 PM - Ruben S. Montero

feature #3782: Update quotas for snapshots on non-persistent disks on disk detach and vm done

Revision 9964cb27 - 06/17/2015 03:45 PM - Ruben S. Montero

feature #3782: Quotas for single snapshot deletion

Revision a2e49d42 - 06/18/2015 10:30 AM - Ruben S. Montero

feature #3782: Snapshots do not increment IMAGE counter in datastore for persistent images

Revision f9bec6a0 - 06/19/2015 10:49 AM - Jaime Melis

feature #3782: Return disk and snapshot information for Ceph disks

Revision dfdbc516 - 06/19/2015 11:47 AM - Ruben S. Montero

feature #3782: Quotas for one.image snapshots actions

Revision d310b310 - 06/22/2015 05:17 PM - Javi Fontan

feature #3782: snap_create for persistent qcow2

Revision 96bea2e6 - 06/23/2015 01:20 PM - Javi Fontan

feature #3782: cpds support for qcow2 snapshots

Revision c6d7c604 - 06/24/2015 12:19 PM - Ruben S. Montero

feature #3782: Make IMAGE attribute mandatory

Revision 4df224e5 - 06/24/2015 01:56 PM - Jaime Melis

feature #3782: Add prefix to attach disk CLI options

Revision ce791d1a - 06/24/2015 03:24 PM - Javi Fontan

feature #3782: tm/snap_revert for persistent qcow2

Revision 96e340de - 06/24/2015 03:24 PM - Javi Fontan

feature #3782: remove duplicate code in qcow2 snap_create

Revision 59f16fd4 - 06/24/2015 09:17 PM - Javi Fontan

feature #3782: ds/snap_revert for qcow2

Revision a1f536dc - 06/24/2015 09:29 PM - Javi Fontan

feature #3782: ds/snap_delete for qcow2

Revision dad5461f - 06/24/2015 10:00 PM - Javi Fontan

feature #3782: ds/snap_flatten for qcow2

Revision 1fd7c533 - 06/24/2015 11:04 PM - Javi Fontan

feature #3782: tm/snap_delete for persistent qcow2

Revision 803592d4 - 06/25/2015 01:36 PM - Ruben S. Montero

feature #3782: Snapshots in SUSPENDED state + dummy changed to exit 0

Revision 87aec8c9 - 06/25/2015 08:22 PM - Javi Fontan

feature #3782: make tm qcow2 snap commands use ssh

Revision 106e8055 - 06/26/2015 08:22 PM - Javi Fontan

feature #3782: Add snap actions to tm/shared

Revision a5e64a37 - 06/29/2015 10:35 AM - Carlos Martín

Feature #3782: Add disk snapshots to sunstone

Revision e00ef94e - 06/29/2015 10:42 AM - Carlos Martín

Feature #3782: Update Sunstone minimized files

Revision cc686d39 - 06/29/2015 02:52 PM - Carlos Martín

Feature #3782: Implement disk-saveas in sunstone

Revision bf5273a4 - 06/29/2015 02:57 PM - Carlos Martín

Feature #3782: Update Sunstone minimized files

Revision c9fc7a36 - 06/30/2015 02:15 PM - Carlos Martín

Feature #3782: Add image snapshots to sunstone

Revision d728d0f9 - 06/30/2015 02:53 PM - Carlos Martín

Feature #3782: Add image snapshots to java api

Revision 6ca1ff4b - 06/30/2015 06:45 PM - Ruben S. Montero

feature #3782: Delete snapshots in RUNNING. Added new states for disk snapshots in running

Revision 282a813d - 06/30/2015 06:45 PM - Ruben S. Montero

feature #3782: Removed unused method

Revision 26da6441 - 06/30/2015 09:36 PM - Javi Fontan

feature #3782: add snapshot functionality to tm/ssh

Revision 10442f81 - 07/01/2015 11:38 AM - Ruben S. Montero

feature #3782: Snapshots in RUNNING state

Revision 5a2d65c5 - 07/01/2015 12:43 PM - Ruben S. Montero

feature #3782: Restore line removed accidentally

Revision e183a8c2 - 07/01/2015 01:13 PM - Ruben S. Montero

feature #3782: Revert in RUNNING state

Revision 143017e3 - 07/01/2015 01:43 PM - Carlos Martín

Feature #3782: Minor fixes to oned code

Revision 201fc040 - 07/01/2015 03:42 PM - Ruben S. Montero

feature #3782: Updated not_supported action to make use of logging functions

Revision 1411e4ad - 07/01/2015 03:57 PM - Carlos Martín

Feature #3782: Use int instead of unsigned int for snapshot IDs

Revision d4c1170e - 07/01/2015 03:57 PM - Carlos Martín

Feature #3782: Remove unused param from ImageManager::set_image_snapshots

Revision fc831b7c - 07/01/2015 04:34 PM - Carlos Martín

Feature #3782: Add disk snapshot actions to java api

Revision 7972b866 - 07/02/2015 10:33 AM - Carlos Martín

Feature #3782: Add confirmation dialogs to snapshot actions

Revision 9e24ff5a - 07/06/2015 03:39 PM - Carlos Martín

Feature #3782: Add image snapshot attributes in onedb upgrade

Revision da86175e - 07/07/2015 02:22 PM - Carlos Martín

Feature #3782: Allow to revert to active disk snapshot in sunstone

Revision 3c3050b1 - 07/07/2015 04:46 PM - Carlos Martín

Feature #3782: Add disk-snapshot commands to states-complete.dot

Revision 8a00683f - 07/07/2015 04:46 PM - Carlos Martín

Feature #3782: Add disk-snapshot commands to sunstone state-actions.js

Revision ae86849b - 07/08/2015 03:58 PM - Ruben S. Montero

feature #3782: Snapshots (create & revert operations) in running can be performed either in a suspend/snap/restore cycle or in a detach/snap/attach one. Allow to revert to the current snapshot

Revision 27e3ac28 - 07/09/2015 11:13 AM - Carlos Martín

Feature #3782: Rename disk snapshot tag to name

Revision 0ad70f21 - 07/14/2015 03:53 PM - Ruben S. Montero

feature #3718: Fix errors in acct with new monitor. feature #3782: Fix acct to account for new snapshots in running. feature #3718: Reset CPU and MEMORY counters

Revision 8a6bc86f - 07/15/2015 09:39 AM - Carlos Martín

Feature #3782: Add new history reasons to xsd and sunstone

Revision 5004376e - 07/16/2015 04:05 PM - Carlos Martín

Feature #3782: Update fsck for new DS and volatile quotas

Revision ecb8d5d5 - 07/17/2015 04:35 PM - Carlos Martín

Feature #3782: Disk-snapshot actions operate on persistent images

For persistent images, quota usage is taken from the image owner.

The image snapshots are now in sync during the VM lifecycle, not just when the VM is shutdown. This was needed because the quota size is now the image size + image quotas, and it caused problems if the owner was changed with snapshots pending to be saved back to the image.

Revision 5213c5a7 - 07/20/2015 10:58 AM - Carlos Martín

Feature #3782: Update deprecated alias for saveas

Revision c9192651 - 07/20/2015 02:00 PM - Carlos Martín

Feature #3782: Update fsck. Image snapshots are in sync with their VM disks

Revision ca60bcb9 - 07/20/2015 04:35 PM - Jaime Melis

Feature #3782: Add ATTACH=YES xml to DISK, not TARGET

Revision 35a5f247 - 07/22/2015 06:31 PM - Ruben S. Montero

feature #3782: Lazy auth for Snapshots methods

Revision 586686f6 - 07/30/2015 10:34 PM - Javi Fontan

feature #3782: Add --live_snapshots option to vmm

This option make vmm skip detach/suspend phase for live snapshots. By default it does live snapshots when vmm is "kvm" and tm is "qcow2"

Revision 77a72e9d - 07/30/2015 10:37 PM - Javi Fontan

feature #3782: support multiple xpath values

xpath.rb will return multiple values if the xpath starts with "%m%"

Revision 5af94991 - 07/30/2015 10:39 PM - Javi Fontan

feature #3782: live snapshots for qcow2

Revision c387ea23 - 07/31/2015 04:40 PM - Javi Fontan

feature #3782: Fix merge with resize clone

Revision 95de9ad5 - 08/06/2015 03:17 PM - Carlos Martín

Feature #3782: Remove disk-snapshot action for volatile disks

Revision 4b4dc90e - 08/06/2015 05:33 PM - Jaime Melis

Feature #3782: Alternate strategy for qcow2/snap_create

Revision a9367b01 - 08/10/2015 10:12 AM - Jaime Melis

Revert "Feature #3782: Alternate strategy for qcow2/snap_create"

This reverts commit 4b4dc90eb62b2cfff15a98576cb7c40cfaa9c90.

Revision b200e9f9 - 08/11/2015 10:56 AM - Jaime Melis

Feature #3782: Clean code, refactor and fix concept errors

Revision 7e7e30cb - 08/11/2015 11:15 AM - Jaime Melis

Feature #3782: qcow2/ln is no longer a symlink to shared/ln

Revision 5fc420b5 - 08/11/2015 02:48 PM - Jaime Melis

Feature #3782: fs ds drivers now handle qcow2 and

raw in the snap_* actions

Revision 40086e5c - 08/13/2015 01:19 PM - Jaime Melis

Feature #3782: Refactor and separate snap_create into two:
snap_create and snap_create_live

Revision 3ba03db8 - 08/13/2015 01:38 PM - Jaime Melis

Feature #3782: make 'suspend' the default disk snapshot technique

Revision a4b60a16 - 08/14/2015 10:33 AM - Jaime Melis

Feature #3782: disable default live disk snapshot

Revision 838b68c6 - 08/24/2015 01:02 PM - Javi Fontan

feature #3782: bug in qcow2 cpds

Revision ed0cfeb6 - 08/24/2015 01:05 PM - Javi Fontan

feature #3782: bugs in qcow2 snap_create_live

Revision d800652d - 08/24/2015 01:40 PM - Javi Fontan

feature #3782: create snap dir link for persistent images

Revision d840ce3e - 08/26/2015 02:05 PM - Javi Fontan

feature #3782: fix monitoring of persistent images

Revision 0a715ca0 - 09/04/2015 12:38 PM - Jaime Melis

Feature #3782: Display correct states for disk-snapshot-* operations

History

#1 - 05/11/2015 05:14 PM - Ruben S. Montero

- Tracker changed from Request to Feature
- Category set to Core & System
- Priority changed from Normal to Sponsored
- Target version set to Release 4.14

#2 - 05/13/2015 04:19 PM - Ruben S. Montero

- Assignee set to Ruben S. Montero

#3 - 05/28/2015 04:47 PM - Gerben Meijer

I've reviewed the commits and as far as I can see, the current iteration of this only allows snapshots to be made with the VM powered off?

I think it would be beneficial to include the ability to snapshot without powering off the VM. I am well aware of the implications of snapshotting a disk live outside the OS when there is no way to know if the data is flushed to disk, however there are a few good reasons why this is still something that is good to have:

- A live disk snapshot in an unknown state is better than no snapshot at all, specifically for cases where a VM cannot be powered off due to customer requirements
- In most cases a live snapshot of a disk for a VM using a journalled filesystem without a guest agent running is recoverable without issues (provided that there is no write back caching going on in qemu, which may not apply to CEPH but does apply to other potential storage systems that support snapshots)
- The fsfreeze feature in the guest agent can be leveraged to trigger a freeze in the guest to ensure sync status, then take a snapshot outside of the VM, and finally thaw it after the snapshot action has been taken (see <http://wiki.qemu.org/Features/QAPI/GuestAgent> for details)

I'm not aware of guest agent support, but that can be implemented as a new feature at a later date - live snapshotting does not explicitly require it and documentation can reflect the risks, with a safe default of not allowing it if so desired.

#4 - 05/28/2015 04:57 PM - Ruben S. Montero

Gerben Meijer wrote:

I've reviewed the commits and as far as I can see, the current iteration of this only allows snapshots to be made with the VM powered off?

I think it would be beneficial to include the ability to snapshot without powering off the VM. I am well aware of the implications of snapshotting a disk live outside the OS when there is no way to know if the data is flushed to disk, however there are a few good reasons why this is still something that is good to have:

- A live disk snapshot in an unknown state is better than no snapshot at all, specifically for cases where a VM cannot be powered off due to customer requirements*
- In most cases a live snapshot of a disk for a VM using a journalled filesystem without a guest agent running is recoverable without issues (provided that there is no write back caching going on in qemu, which may not apply to CEPH but does apply to other potential storage systems that support snapshots)*
- The fsfreeze feature in the guest agent can be leveraged to trigger a freeze in the guest to ensure sync status, then take a snapshot outside of the VM, and finally thaw it after the snapshot action has been taken (see <http://wiki.qemu.org/Features/QAPI/GuestAgent> for details)*

I'm not aware of guest agent support, but that can be implemented as a new feature at a later date - live snapshotting does not explicitly require it and documentation can reflect the risks, with a safe default of not allowing it if so desired.

Hi Gerben Meijer

Yes you are right, this is the initial code to test things, you know ;). Once most of the machinery is in place, we'll add the RUNNING state as possible and test things. For example, we want to stress libvirt when reverting to a snapshot (the original block or file will be replaced by another one and that may not play nicely with qemu/libvirt)...

THANKS FOR YOUR FEEDBACK :)

#5 - 06/08/2015 10:05 PM - Ruben S. Montero

Needed modifications:

Compatibility

- Clone cannot be perform on images with snapshots
- Image with snapshots can be only of OS and DATA type (chtype enforce that)
- Image with snapshots can be only persistent (persistent api enforce that)

Upgrade

- 3 new states for disk-snapshot operations
- New field <TARGET_SNAPSHOT> that marks active snapshot operation (similar to CLONING_ID)
- New driver operations

- *TM: snap_create, snap_delete, snap_revert + update of mvds and rm*
- *DS: snap_delete, snap_revert, snap_flatten + update of delete*

- Migrator to add IMAGE element

API

- disk and image operations. OCA JAVA missing

Sunstone

- Refactor of VM disk tab
- Add snapshot tab for Images

#6 - 06/10/2015 06:29 PM - Ruben S. Montero

Updates:

- level JAVA API with the new one.vm.disksaveas, this replaces the old disk_snapshot api call, and onevm.saveas method
- Check that JSON classes used by Sunstone are properly updated
- Sunstone needs to re-work the snapshot/saveas interface
- Datamodel, removed DISK fields SAVE_AS_SOURCE, SAVE_AS and SAVE_AS_ACTIVE are not in use anymore, code should not check against those.
- Check EC2 API's

Compatibility

- Applications using disk_snapshot API call needs to upgrade to the new functionality
- New API call, with optional snap_id parameter
- Deferred snapshots are superseded by the new snapshotting functionality
- CPDS has a new argument, positional parsing of arguments by custom scripts needs to be updated
- Removed cancelsaveas, not needed anymore

Upgrade guide

- VMs with snapshot deferred will not perform the saveas operation, needs to be manually triggered, and placeholder images maybe left in locked state (need to remove them)*

#7 - 06/16/2015 12:05 AM - Ruben S. Montero

Evaluate migrators for VMs with SAVE_AS=YES (deferred)

#8 - 06/23/2015 03:18 PM - Jaime Melis

libvirt/qemu currently does not allow reopening the disk to account for a disk change, which is what's happening in the snapshot-create operation. For the moment we will allow snapshot-create to be triggered in POWEROFF and SUSPENDED states. This will change once libvirt exposes an interface

to make this change.

#9 - 07/15/2015 12:42 PM - Stefan Kooman

@Jaime Meles libvirt/qemu currently does not allow reopening the disk to account for a disk change, which is what's happening in the snapshot-create operation. For the moment we will allow snapshot-create to be triggered in POWEROFF and SUSPENDED states. This will change once libvirt exposes an interface to make this change.

^^ when snapshots are created using "virsh snapshot-create-as" there is no need to notify libvirt/qemu of a disk change, as this is done automatically (intrinsic action of snapshot-create-as operation").

virsh help snapshot-create-as

NAME

snapshot-create-as - Create a snapshot from a set of args

SYNOPSIS

snapshot-create-as <domain> [--name <string>] [--description <string>] [--print-xml] [--no-metadata] [--halt] [--disk-only] [--reuse-external] [--quiesce] [--atomic] [--live] [--memspec <string>] [--diskspec <string>]..

According to this mail thread: <https://www.redhat.com/archives/libvirt-users/2013-September/msg00146.html>, this feature request: <http://tracker.ceph.com/issues/2773> and the above output it *should* be possible to create a "--disk-only" snapshot on a live running VM. I have tested this on a qcow2 backed VM and this works. Unfortunately I do not (yet) have a CEPH cluster to test against but it's worth a try.

--diskspec <string> disk attributes: disk[,snapshot=type][,driver=type][,file=name].

^^ CEPH specific parameters should be filled in.

#10 - 08/17/2015 07:52 AM - Ruben S. Montero

- Status changed from Pending to Closed

- Resolution set to fixed

#11 - 01/18/2016 03:11 AM - Guillaume Michaud

Hi,

Disk snapshot on a live running VM with Ceph backend does not seem to trigger the guest-fsfreeze-freeze/thaw commands on the guest. Is it supported in one-4.14.2?

Followed this procedure : <http://www.sebastien-han.fr/blog/2015/02/09/openstack-perform-consistent-snapshots-with-qemu-guest-agent/>

Works well if I trigger the command manually from the host (/tmp/freeze is filled with I'm frozen/I'm thaw). But nothing happens if I snapshot the disk manually from Sunstone.

#12 - 01/18/2016 04:26 PM - Ruben S. Montero

Hi,

It's been called here

https://github.com/OpenNebula/one/blob/release-4.14.2/src/tm_mad/ceph/snap_create_live#L95

Could it be some problems with libvirt/qemu version?

Guillaume Michaud wrote:

Hi,

Disk snapshot on a live running VM with Ceph backend does not seems to trigger the guest-fsfreeze-freeze/thaw commands on the guest. Is it supported in one-4.14.2?

Followed this procedure : <http://www.sebastien-han.fr/blog/2015/02/09/openstack-perform-consistent-snapshots-with-qemu-guest-agent/>

Works well if I trigger the command manually from the host (/tmp/freeze is filled with I'm frozen/I'm thaw). But nothing happen if I snapshot the disk manually from Sunstone.

#13 - 01/18/2016 05:27 PM - Guillaume Michaud

Hi Ruben,

As I understand, the virsh command is called from the frontend but it is not installed my CentOS 7 frontend. I have to install the libvirt package to get it. I installed my frontend using the official ONE repo and "yum install opennebula-server opennebula-sunstone" command.

Thanks.

Ruben S. Montero wrote:

Hi,

It's been called here

https://github.com/OpenNebula/one/blob/release-4.14.2/src/tm_mad/ceph/snap_create_live#L95

Could it be some problems with libvirt/qemu version?

Guillaume Michaud wrote:

Hi,

Disk snapshot on a live running VM with Ceph backend does not seems to trigger the guest-fsfreeze-freeze/thaw commands on the guest. Is it supported in one-4.14.2?

Followed this procedure : <http://www.sebastien-han.fr/blog/2015/02/09/openstack-perform-consistent-snapshots-with-qemu-guest-agent/>

Works well if I trigger the command manually from the host (/tmp/freeze is filled with I'm frozen/I'm thaw). But nothing happen if I snapshot the disk manually from Sunstone.

#14 - 01/18/2016 07:09 PM - Guillaume Michaud

Installing libvirt-client on the frontend does not make a difference. So I understand that the virsh command is called directly on the host.

Executing "virsh -c qemu:///system domfsfreeze one-546" on the host trigger the fshook correctly in the guest.

Just want to be sure, I'm talking about disk-snapshoting and not VM-snapshoting. On the Storage tab, I'm using the snapshot option for the vda disk (ID-0).

Thanks.