

OpenNebula - Bug #2067

sunstone and rbd

05/17/2013 05:54 PM - George Kissandrakis

Status:	Closed	Start date:	05/17/2013
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:	Sunstone	Estimated time:	0.00 hour
Target version:		Pull request:	
Resolution:	worksforme		
Affected Versions:	OpenNebula 4.0		

Description

I have installed opennebula within sunstone and rbd (ceph)

I have one problem

I create an image in rbd datastore (through sunstone or command line, no difference)

- Case 1 command line

I create a VM from command line

```
onevm create --name vm1 --cpu 1 --mem 1024 --disk image1 --vnc --nic 1
```

the VM uses the image (image1) for disk which is in rbd

Live migrate works perfect

Failover works perfect

the command line for this vm is

```
/usr/bin/kvm -name one-36 [blah...blah] -drive
```

```
file=rbd:one/one-9:auth_supported=none,if=none,id=drive-ide0-0-0,format=raw,cache=none [blah...blah]
```

So the disk is the rbd block device and only that, which is distributed to all physical hosts through ceph

- Case 2 through sunstone

I create a VM template using the same image (image1) (persistent or not)

I create a VM with that template

and the command line is

```
/usr/bin/kvm -name one-35 [blah...blah] -drive
```

```
file=/var/lib/one//datastores/0/35/disk.1,if=none,media=cdrom,id=drive-ide0-0-0,readonly=on,format=raw [blah...blah] -drive
```

```
file=rbd:one/one-8:auth_supported=none,if=none,id=drive-ide0-0-1,format=raw,cache=none [blah...blah]
```

As you see in the second case there is a file as drive, so when i try to migrate or to failover by crashing the server, the VM does not migrate/failover because is missing the file

If i do a shared storage, it migrates the VM but it does not failover because it tries to copy the file from the crashed server(?) and even if it migrates it will loose all data which are in file

So is there a way to create a VM which reads and writes in the rbd directly and not on a file with sunstone?

Is it a bug or lack of feature?

History

#1 - 05/21/2013 10:25 AM - George Kissandrakis

Not a bug,

Just configuration

Any context in template must be unchecked

#2 - 05/28/2013 10:18 AM - Tino Vázquez

Could you please elaborate more on what you mean about "Any context in template must be unchecked"?

#3 - 06/06/2013 11:11 AM - Jaime Melis

- *Status changed from Pending to Closed*

- *Resolution set to worksforme*

Yes George you are right. The problem is the context that's being generated as a file. You can one of two things:

1. Configure the system datastore as shared storage (cephfs for instance! or nfs, etc...)
2. If you're using only ceph based images, you could enhance the ceph drivers and use them as a system datastore. You would need to implement other driver actions, like CONTEXT. More info here: <http://opennebula.org/documentation:rel4.0:sd>