

Open Infrastructure - Task #6872

Task # 6869 (Rejected): Challenge OpenStack and OpenNebula with ucloud

Create ucloud-vm

06/23/2019 05:11 PM - Nico Schottelius

Status:	Closed	Start date:	06/23/2019
Priority:	Normal	Due date:	06/27/2019
Assignee:	Ahmed Bilal	% Done:	60%
Category:		Estimated time:	0.00 hour
Target version:			
PM Check date:	06/26/2019		

Description

This service will run on EVERY vm host and will do the heavy lifting / creating of VMs

For v1, implement the following:

- Watch etcd on key prefix /v1/vm/
- If there is a change, check whether
 - an existing VM on the current host should be killed
 - a new VM should be started
 - On event emit an email to sre@ (mail should be injected by env) -- just initially, later this will be disabled
- Create a VM
 - if there is a new json entry in /v1/vm with a VM that is designated for this host and has not been started
 - Use qemu ... for it (sample below)
- Delete a VM
 - If the status json entry of a VM is TO_STOP, stop it
 - After stopping was successful, put it into STOP
 - We need to register the runtime somewhere (v2)
- Monitor a VM
 - If it is still running
 - If it is not running anymore, notify sysadmins
 - Write to /v1/log/
 - And/or send email to sre@
 - And restart the VM

For v2 implement:

- Every 5 minutes poll etcd
- Check whether the VMs that are supposed to run on this host are running
 - If one missing: send mail to sre@
- Check that not OTHER VMs are running on this host
 - If there is one that doesn't belong here: send mail to sre@

Sample qemu commands from production

- This is way too long, we can trim it down.
- We will also use the -name attribute
- We will also use ceph and smp
- We will also use uuid
- We will run VNC
 - v1: tcp
 - v4: unix sockets

```
oneadmin 3641289 0.5 0.1 2577776 407572 ? S1 Jun21 59:10 /usr/bin/qemu-system-x86_64 -name guest=one-24442,debug-threads=on -S -object secret,id=masterKey0,format=raw,file=/var/lib/libvirt/qemu/domain-24-one-24442/master-key.aes -machine pc-i440fx-2.8,accel=kvm,usb=off,dump-guest-core=off -m 1024 -realtime mlock=off -smp 1,sockets=1,cores=1,threads=1 -uuid d902efb7-2c66-4685-8914-95d38a717c90 -no-user-config -nodefaults -chardev socket,id=charmonitor,path=/var/lib/libvirt/qemu/domain-24-one-24442/monitor.sock,server,nowait -mon chardev=charmonitor,id=monitor,mode=control -rtc base=utc -no-shutdown -boot strict=on -device piix3-usb-uhci,id=usb,bus=pci.0,addr=0x1.0x2 -object secret,id=virtio-disk0-secret0,data=PR218Q61gRpXffWtRuIqQntokENHbAaBJGDDG66u3lg=,keyid=maste
```

```
rKey0,iv=XUHyWqKdHWGuva4eX+5qPw==,format=base64 -drive file=rbd:one/one-449-24442-0:id=libvirt:auth_supported=cephx\;none:mon_host=ceph1\:6789\;ceph2\:6789\;ceph3\:6789,file.password-secret=virtio-disk0-secret0,format=raw,if=none,id=drive-virtio-disk0,cache=none -device virtio-blk-pci,scsi=off,bus=pci.0,addr=0x4,drive=drive-virtio-disk0,id=virtio-disk0,bootindex=1 -drive file=/var/lib/one/datastores/100/24442/disk.1,format=raw,if=none,id=drive-ide0-0-0,readonly=on -device ide-cd,bus=ide.0,unit=0,drive=drive-ide0-0-0,id=ide0-0-0 -netdev tap,fd=35,id=hostnet0,vhost=on,vhostfd=37 -device virtio-net-pci,netdev=hostnet0,id=net0,mac=02:00:b3:39:79:75,bus=pci.0,addr=0x3 -vnc [::]:4162 -device cirrus-vga,id=video0,bus=pci.0,addr=0x2 -device virtio-balloon-pci,id=balloon0,bus=pci.0,addr=0x5 -msg timestamp=on
```

History

#1 - 06/26/2019 07:36 AM - William Colmenares

- % Done changed from 0 to 60
- Due date set to 06/27/2019

Added email on events, working in V2

#2 - 06/26/2019 04:28 PM - Jason Kim

- PM Check date set to 06/26/2019

#3 - 06/26/2019 05:44 PM - William Colmenares

- Status changed from New to In Progress

Working on v2, finding a way to display the vm ID on kill request

#4 - 06/29/2019 09:59 AM - Nico Schottelius

- Description updated
- Status changed from In Progress to New
- Assignee changed from William Colmenares to Ahmed Bilal

#5 - 06/30/2019 02:23 PM - Ahmed Bilal

- Instead of **TO_STOP** I think we have to react on **REQUESTED_DELETE** as mentioned [here](#)
- Similarly, instead of **TO_STOP** --> **STOP** we need **REQUESTED_DELETE** --> **DELETED**
- For VMs, which are killed in action (abnormally by some bug/error/user operation) we would set its status to "KILLED"

#6 - 06/30/2019 05:08 PM - Ahmed Bilal

- Status changed from New to In Progress

#7 - 07/01/2019 06:00 PM - Nico Schottelius

I think we need to distinguish between DELETE and SUSPEND where

DELETE = VM is irrecoverable dead

SUSPEND = VM can be respawn and keeps IP and image

redmine@ungleich.ch writes:

#8 - 07/04/2019 02:18 PM - Ahmed Bilal

- Status changed from In Progress to Feedback

#9 - 07/15/2019 04:30 PM - Ahmed Bilal

- Status changed from Feedback to In Progress

#10 - 11/28/2019 08:33 PM - Ahmed Bilal

- Status changed from In Progress to Closed